arXiv:1403.3109v1 [cs.IT] 12 Mar 2014

Sparse Recovery with Linear and Nonlinear Observations: Dependent and Noisy Data

Cem Aksoylar and Venkatesh Saligrama

Abstract

We formulate sparse support recovery as a salient set identification problem and use information-theoretic analyses to characterize the recovery performance and sample complexity. We consider a very general model where we are not restricted to linear models or specific distributions. We state non-asymptotic bounds on recovery probability and a tight mutual information formula for sample complexity. We evaluate our bounds for applications such as sparse linear regression and explicitly characterize effects of correlation or noisy features on recovery performance. We show improvements upon previous work and identify gaps between the performance of recovery algorithms and fundamental information.

1 Introduction

In this paper, we consider problems where among a set of D variables/features $X = (X_1, \ldots, X_D)$, only K variables (indexed by set S) are directly relevant to the observation/label Y. These types of problems frequently arise in a number of scenarios in high-dimensional analysis, such as compressive sensing [1], feature selection in learning [2] or other high-dimensional problems with an inherent low-dimensional structure. We formulate these problems with the following Markovian property: Given $X_S = \{X_k\}_{k \in S}$, observation Y is independent of other variables $\{X_k\}_{k \notin S}$, i.e.,

$$P(Y|X) = P(Y|X_S). \tag{1}$$

Given N sample pairs $(X^N, Y^N) = \{(X^{(1)}, Y^{(1)}), \ldots, (X^{(N)}, Y^{(N)})\}$, our goal is to identify the set of relevant/salient variables S from these N samples. Our analysis aims to characterize the recovery performance probabilistically and establish necessary & sufficient conditions on N in order to recover the set S with an arbitrarily small error probability in terms of K, D and model parameters such as the signal-to-noise ratio (SNR).

As an illustrative example, consider the sparse linear regression model, given by $Y^N = X^N\beta + W^N$, where S is the support of sparse random vector β and random noise W^N assumed to be independent of X^N and β. The elements in a row of the signal matrix correspond to variables $X_1, \ldots, X_D$. Each row is a realization of X and X^N is formed from sampled rows. Markov assumption (1) is satisfied, since each $Y^{(n)}$ depends only on the linear combination of the elements $X_S^{(n)}$ that correspond to the support of β. The coefficients of this combination are given by β_S, viewed as a random "nuisance" parameter in the observation model. This perspective also holds for non-linear models, thus unifying many sparse recovery problems.

Information-theoretic approaches with relation to channel coding [3] have been considered in previous work for different application areas, where the salient set S is seen as a message encoded by X^N and is recovered from outputs Y^N. Specifically, the problem of group testing was formulated in a similar framework in Russian literature [4–8] and in [9]. [10] has followed a similar approach to [9] to obtain sample complexity results for general sparse signal processing models. Note that the identification problem formulated here has key differences with channel coding, namely the inability to "code" the variables X^N and different messages/sets overlapping and thus sharing codewords.

Previous work on general models, namely [10] has severe limitations related to the specific analysis techniques used. In this work we overcome several of these limitations, where we (1) consider correlated

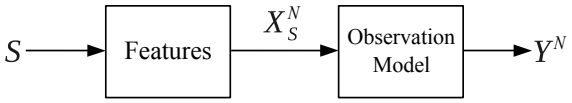

Figure 1: Channel model.

variables instead of independent and identically distributed (IID) variables and relate correlation to sample complexity, (2) present a non-asymptotic analysis through probability of recovery bounds instead of solely asymptotic analyses for sample complexity and (3) state more general results for the high-dimensional sparsity regime where sparsity scales with the number of variables.

The bounds we present for the sample complexity are of the form

$$N\, I(X_S; Y) > \log \binom{D}{K}, \tag{2}$$

which can be interpreted as follows: The right side of the inequality is the number of bits required to represent all sets S of size K. On the left side, the mutual information term [3] represents the uncertainty reduction on the output Y when given the input X_S, in bits per sample. This term essentially quantifies the "capacity" of the observation model $P(Y|X_S)$. (2) is then a statement that uncertainty reduction with N samples should exceed uncertainty of set S.

Furthermore, our analysis also provides us with a sharp exponential upper bound on the probability of error in identifying the salient set. This bound can be computed easily and a closed form expression can be obtained for some applications, such as in the case of linear models. We compute these bounds for the sparse linear model described, where we explicitly characterize the effects of correlation, SNR or noisy variables.

Many models sharing the common structure of sparsity satisfy Markovian assumption (1). These include sparse linear regression or compressive sensing (CS) [1], probit regression or 1-bit CS [11,12], group testing [9], sparse logistic regression and multiple regression models [13] with group sparsity property. In addition, variants of these problems can be considered, e.g., with noisy or missing data where variables are not fully observed (see [14]).

1.1 Related Work and Contributions

The problem of sparse recovery and in particular information-theoretic (IT) analysis is extensive. We only describe work closely related to this paper. Much of the IT literature deals with linear models and mean-squared estimation of β in the sparse linear model with sub-Gaussian assumptions on variables X^N. Below we list the contributions of our approach and contrast it with some of the related work in the literature. The Markovian problem formulation follows along the lines of [10, 15], which we repeat here for the sake of exposition.

Unifying framework through Markovianity: Much of the literature on sparse recovery is specialized with tailored algorithms for different problems. For instance, lasso for linear regression [16,17], relaxed integer programs for group testing [18], convex programs for 1-bit quantization [12], projected gradient descent for sparse regression with noisy and missing data [14] and other general forms of penalization. While all of these problems share an underlying sparse structure, it is conceptually unclear from a purely IT perspective, how they come together from an inference point-of-view. Our Markovian viewpoint of (1) unifies these different sparse problems from an inference perspective.

Discrete objects with continuous observations: While [17, 19–24] describe IT bounds for sparsity pattern recovery to recover S, they exclusively focus on the linear sub-Gaussian setting. Furthermore, their approach is *circuitous*. Indeed, they rely on first estimating the sparse vector β, which is then thresholded to obtain S. This not only complicates the analysis and introduces unnecessary assumptions on β but also obfuscates the distinction between signal estimation vs. support discovery. It is well-known that if support is known, signal estimation is easy and least-squares estimates are reliable. At a conceptual level IT tools such as Fano's inequality and capacity theorems are powerful tools for inferring about discrete objects (messages)

given continuous observations. Indeed, to exploit IT tools, [19–24] resort to one of the following strategies: (a) Use IT tools only for necessity part by assuming a special case of discrete β and derive sufficiency with some well-known algorithm (lasso, basis pursuit etc.); or (b) find a ϵ-cover for β in some metric space (which requires imposing extra assumptions) and reduce β to a discrete object. In contrast our approach lifts these assumptions and focuses on the natural discrete object S. Our result shows that indeed the discrete part, namely, uncertainty support pattern is the dominating factor and not β itself.

Furthermore, prior work relied heavily on the design of sampling matrices with special structures such as Gaussian ensembles and RIP matrices, which is a key difference from the setting we consider herein as for our purpose we do not always have the freedom to design the matrix X. We do not make explicit assumptions about the structure of the sensing matrix, such as the restricted isometry property [25] or incoherence properties [16], or about the distribution of the matrix elements, such as sub-Gaussianity. Also, the existing IT bounds, which are largely based on Gaussian ensembles, are limited to the linear CS model, and hence not suitable for the non-linear models we consider herein.

Information-theoretic tight error bounds: Through our analysis of the ML decoder, we obtain a tight upper bound on the probability of error of support recovery, in addition to necessary and sufficient conditions on the sample complexity. We compute this upper bound explicitly for popular problems such as sparse linear regression and its variants. We compare the information-theoretic bound to the performance of practical algorithms used to solve the sparse recovery problem, such as lasso [16, 17] or orthogonal matching pursuit (OMP) variants [26] and illustrate large gaps between their performance and our bounds. The presence of these gaps show that there is still room to improve the performance of practical algorithms for solving support recovery problems.

Bounds for new sparse recovery problems: Our unifying approach also allows the study of problems that are not previously analyzed, or that are not easily analyzed using other approaches. These types of problems may include sparse recovery with novel observation models, or existing models with different distributions of variables or noise. Due to our Markovian formulation, obtaining necessary and sufficient conditions and error bounds only necessitate computation of simple mutual information expressions and an error exponent expression.

As mentioned in the introduction, the identification problem was formulated in a channel coding framework in [9] and in the Russian literature [4–8]. This analysis was extended to general sparse signal processing models with IID variables and latent variable observation model in [10] and [15].

We consider the analysis of models with correlated variables, specifically conditionally IID variables X given a latent parameter θ. We also state a non-asymptotic bound on the probability of error, which in turn allows us to identify performance gaps between practical algorithms and our information-theoretic results. In addition, we consider a general scaling regime where $K = O(D)$ for linear models and variants through this bound. We also introduce the noisy data framework and explicitly characterize the recovery performance w.r.t. the noise variance.

1.2 Problem setup

Notation. We represent variables with row vectors and samples as different rows to obtain a $N \times D$ variable matrix, while the observation samples form a column vector. In that context, subscripts are used for column indexing and superscripts with parentheses are used for row indexing in vectors and matrices. log is used to denote logarithm to the base 2.

Problem setup. We observe the realizations of N variable-observation pairs (X^N, Y^N) with each sample $(X^{(n)}, Y^{(n)})$, $n = 1, 2, \ldots, N$. Observations Y are given by $P(Y|X_S, \beta_S)$ with latent model parameter $\beta_S \sim P(\beta_S)$ and satisfy the Markovian property (1), where $|S| \leq K$ with known $K \ll D$. Observation parameters β_S correspond to the coefficients on the support of the sparse vector in sparse recovery problems. For simplicity of exposition we consider the case $|S| = K$. The variables $X^{(n)}$ are IID across $n = 1, \ldots, N$. However, the observations $Y^{(n)}$ are independent for different n only when conditioned on β_S. Our goal is to identify the set S from the N samples of variables and the associated observations (X^N, Y^N), with an arbitrarily small average error probability.

We index the different sets of size K as S_ω, so that S_ω is a set of K indices corresponding to the ω-th set of variables. Since there are D variables in total, there are $\binom{D}{K}$ such sets, hence $\omega \in \left\{1, 2, \dots \binom{D}{K}\right\}$.

Let $\hat{S}(X^N, Y^N)$ denote the estimate of the set S which is random due to the randomness in X and Y and let $P(E)$ denote the average probability of error, averaged over all sets S of size K, variables X^N and observations Y^N, i.e., $P(E) = \Pr[\hat{S}(X^N, Y^N) \neq S]$.

2 Recovery and Error Bounds

Central to our analysis are the following four assumptions, which we utilize in order to analyze the probability of error in recovering the salient set and to obtain bounds on sample complexity.

(A1) **Equiprobable support:** Any set $S_\omega \subset \{1, \dots, D\}$ with K elements is equally likely *a priori* to be the salient set. We assume no prior knowledge of the salient set S among $\binom{D}{K}$ such sets.

(A2) **Conditional independence:** The observation/label Y is conditionally independent of other variables given X_S, variables with indices in S, i.e., $P(Y|X) = P(Y|X_S)$.

(A3) **Conditionally IID variables:** The variables $X_1, \dots, X_D$ are IID conditioned on a latent parameter θ. We elaborate on this property in the following sections.

(A4) **Observation model symmetry:** For any permutation mapping π, $P(Y|X_S) = P(Y|X_{\pi(S)})$, i.e., the observations are independent of the ordering of variables. This is not a very restrictive assumption since the asymmetry w.r.t. the indices can be incorporated into β_S, as the symmetry is assumed for the observation model averaged over β_S.

With only these four general assumptions, we are able to identify bounds that we state in the next section, for a general class of problems. We now elaborate on some of the assumptions above.

Conditional Independence

A simple example with the conditional independence property is the sparse linear model mentioned in the introduction. In this model every observation is given by the model

$$Y = \langle X, \beta \rangle + W = \langle X_S, \beta_S \rangle + W$$

with noise W, which can also be extended to nonlinear models $Y = f(\langle X_S, \beta_S \rangle + W)$, for a function $f : \mathbb{R} \to \mathbb{R}$.

A different nonlinear example is the group testing model, which is detailed in [9, 15]. In this model observations are Boolean test results, while the variables are Boolean test inclusion indicators for each test and item. Each test result is positive if and only if any item from a certain defective set of items is included in that test, i.e.,

$$Y = \bigvee_{k \in S} X_k.$$

Denoting the set of defective items with S, it is then clear that the conditional independence holds.

Conditionally IID Variables

For conditionally IID variables, the joint distribution of variables can be written as

$$P(X_1, \dots, X_D) = \int_\Theta \prod_{k=1}^{D} P(X_k|\theta) P(\theta) \, \mathrm{d}\theta,$$

where $\theta \in \Theta$ is the latent coupling parameter with density $P(\theta)$. (A3) appears restrictive and so we describe a few examples and extensions to build intuition.

Bouquet model [27] arises in sparsity-based face recognition and given by $X_k = \mu + W_k, \quad k = 1, \ldots, D$, with $W_k \sim \mathcal{N}(0, \sigma_W^2)$ IID across k and $\mu \sim \mathcal{N}(0, \sigma_\mu^2)$. It can be seen that two variables X_i and X_j are dependent and correlated with correlation coefficient $\rho = \sigma_\mu^2/(\sigma_\mu^2 + \sigma_W^2)$ but IID conditioned on μ.

Meta parameters: We can account for several possibilities by selectively introducing meta parameters. For instance, we can let $X_k = \alpha_k^\top \mu$ with $\mu \sim \mathcal{N}(0, I_D)$ and IID random vectors α_k. Here $(\{X_k\} \mid \{\alpha_k\}, \mu)$ are independent, not identically distributed with $E(X_k X_j \mid \{\alpha_k\}, \mu) = \alpha_k^\top \alpha_j$. Nevertheless, our results also extend to this setting. Note that X_k's are exchangeable. Indeed, there is a close connection between conditional IID random variables and exchangeable random variables through de Finetti's theorem [28–30].

2.1 Recovery Conditions and Error Bounds

To derive the upper bound on recovery error and sufficiency bound for the required number of samples, we analyze the error probability of a Maximum Likelihood (ML) decoder [31]. The decoder goes through all $\binom{D}{K}$ sets of size K and chooses the set S_{ω^*} for which observation Y^N is most likely, i.e.,

$$P(Y^N | X^N_{S_{\omega^*}}) > P(Y^N | X^N_{S_\omega}), \quad \forall \omega \neq \omega^*. \tag{3}$$

An error occurs if any set other than the true set S is more likely. This ML decoder is a minimum probability of error decoder assuming uniform prior on the candidate sets of variables. Note that the ML decoder requires the knowledge of the observation model $P(Y|X_S, \beta_S)$ and the prior $P(\beta_S)$. Next, we derive an upper bound on the error probability $P(E)$ of the ML decoder, averaged over all sets, data realizations and observations.

Our methodology for the analysis is as follows: To deal with scenarios where a candidate set S_ω and true set S have overlapping elements (and thus $X^N_{S_\omega}$ and X^N_S share certain columns), we define the error event E_i as the event of mistaking the true set for a set which differs from the true set S in exactly i variables, i.e., there exists some set which differs from the true set in i variables and is more likely to the decoder. Note that $E = \bigcup_{i=1}^K E_i$. Then for each i we use an analysis based on the characterization of error exponents as in [31] to obtain an upper bound on $P(E_i)$, which leads to Theorem 2.1 and a sufficient condition on N. We derive a matching necessity bound on N with an argument based on Fano's inequality [3].

Our first main result is the following theorem, which states a non-asymptotic upper bound on the probability of error of exact support recovery.

Theorem 2.1. *Under the assumptions (A1)-(A4), the probability of error $P(E)$ that a set other than S is selected by the ML decoder is bounded from above by*

$$P(E) \leq \min_{\delta \in [0,1]} \sum_{i=1}^{K} 2^{-N\left(E_o(\delta) - \delta \frac{\log \binom{D-K}{i}\binom{K}{i}}{N}\right)}, \tag{4}$$

where

$$E_o(\delta) = -\frac{1}{N} \log E_{\theta^N} \left[\sum_{Y^N} \sum_{X^N_{\mathcal{S}^2}} P(X^N_{\mathcal{S}^2} | \theta^N) \left(\sum_{X^N_{\mathcal{S}^1}} P(X^N_{\mathcal{S}^1} | \theta^N) P(Y^N | X^N_{\mathcal{S}^1}, X^N_{\mathcal{S}^2})^{\frac{1}{1+\delta}} \right)^{1+\delta} \right],$$

for $0 \leq \delta \leq 1$. $(\mathcal{S}^1, \mathcal{S}^2)$ denotes a disjoint partition of the true set of variables S with cardinalities i and $K - i$, $X^N_{\mathcal{S}^1}$ and $X^N_{\mathcal{S}^2}$ are the corresponding disjoint partitions of the $N \times K$ input $X^N_{\mathcal{S}}$ of sizes $N \times i$ and $N \times (K - i)$, respectively. θ is the parameter in the cond. IID representation. The bound holds for any (N, K, D).

Remark 2.1. *For fixed and known β_S, observations $Y^{(n)}$ are independent and $E_o(\delta)$ simplifies to*

$$E_o(\delta) = -\log E_\theta \left[\sum_Y \sum_{X_{\mathcal{S}^2}} P(X_{\mathcal{S}^2}|\theta) \left(\sum_{X_{\mathcal{S}^1}} P(X_{\mathcal{S}^1}|\theta) P(Y|X_{\mathcal{S}^1}, X_{\mathcal{S}^2})^{\frac{1}{1+\delta}} \right)^{1+\delta} \right].$$

Next, we state our main result for the sample complexity of support recovery. The following theorem provides tight necessary and sufficient conditions on the number of samples N asymptotically for an arbitrarily small average error probability.

Theorem 2.2. *Let $I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2}, \beta_S, \theta)$ be the mutual information between $X_{\mathcal{S}^1}$ and Y conditioned on $X_{\mathcal{S}^2}$, β_S and θ. Under the assumptions (A1)-(A4), a necessary condition on the number of samples N to recover S with an arbitrarily small average probability is given by*

$$N > (1+\epsilon) \max_{i=1,\dots,K} \frac{\log \binom{D-K+i}{i}}{I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2}, \beta_S, \theta)}. \tag{5}$$

Furthermore, if $I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2}, \beta_S, \theta) = \omega(1/\log D)$ for all $i = 1, \dots, K$, (5) is also a sufficient condition. The necessary condition holds for all scalings $K = O(D)$ and $\epsilon = 0$, while the sufficiency bound holds for any fixed K as $D \to \infty$ and $\epsilon > 0$ an arbitrary constant.

Note that the condition that $I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2}, \beta_S, \theta) = \omega(1/\log D)$ is not restrictive, since typically the mutual information per sample depends on the number of salient variables K and not on the total number of variables D and we consider the regime where K is fixed w.r.t. D for the sufficient condition.

IID variables. For IID variables, the mutual information expression in the denominator is $I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2}, \beta_S)$ and further reduces to $I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2})$ for fixed observation parameters β_S.

Interpretation. Intuitively, the condition in (5) can be explained as follows: For each i, the numerator is the number of bits required to represent all sets S_ω with $K-i$ indices known beforehand. The denominator represents the information given by the output variable Y about the remaining i indices $\mathcal{S}^1$, given the subset $\mathcal{S}^2$ of $K-i$ true indices. Hence, the ratio represents the number of samples needed to control i support errors in $\mathcal{S}^1$ and maximization accounts for all possible support errors.

Partial recovery. As we analyze the error probability separately for $i = 1, \dots, K$ support errors in order to obtain the necessity and sufficiency results, it is trivial to determine conditions for *partial* instead of *exact* support recovery. By changing the maximization from over $i = 1, \dots, K$ to $i = \lfloor \alpha K \rfloor, \dots, K$ in (5), the conditions to recover at least $(1-\alpha)K$ of the K support indices can be determined.

Support pattern recovery dominates support coefficient estimation. In the proof of Theorem 2.2, we show that β_S being unknown with prior $P(\beta_S)$ induces a penalty term in the denominator given by $I(\beta_S; X^N_{\mathcal{S}^1}|X^N_{\mathcal{S}^2}, Y^N, \theta^N)/N$, compared to the case where support coefficients β_S are fixed and known. We show that this term is always dominated by $I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2}, \beta_S, \theta)$ provided a mild condition on the mutual information is satisfied, therefore does not affect the sample complexity asymptotically. This shows that recovering support while knowing the support coefficients is as hard as recovering with unknown coefficients, underlying the importance of recovering the support in sparse recovery problems.

3 Applications

In this section, using the result of Theorem 2.1, we provide explicit non-asymptotic upper bounds for the error probability for sparse linear models that may include correlations or noisy variables. We also state asymptotic sample complexity results using the error bounds and Theorem 2.2. We then compare the information-theoretic error bounds we obtained with the recovery performance of practical algorithms.

3.1 Sparse linear regression

We consider the normalized model [21],

$$Y^N = X^N\beta + W^N, \tag{6}$$

where X^N is the $N \times D$ sensing matrix, β is a K-sparse vector of length D with support S and Y^N is the observation vector of length N. We assume $X^{(n)}$ are jointly Gaussian row vectors and IID across rows n. Each element $X_k^{(n)}$ is zero mean and has variance $1/N$. W^N is the IID observation noise, with $W \sim \mathcal{N}\left(0, \frac{1}{\text{SNR}}\right)$. The coefficients of the support, β_S, are either fixed and $|\beta_k| = \sigma$, or IID Gaussian with zero mean and variance σ^2.

Theorem 3.1. *Consider the correlated setup described above. First, SNR $= \Omega(\log D)$ is a necessary condition for recovery. Furthermore, for this SNR we can recover S with average error probability approaching zero if and only if $N = \Omega\left(\frac{K\log(D/K)}{\log(1+(1-\rho)\sigma^2)}\right)$.*

We consider a generalized model, which may include correlations between sensing columns, such that $E[X_k^{(n)}X_{k'}^{(n)}] = \rho/N$. ρ is then the correlation coefficient between two columns. Note that this model is statistically equivalent to the following one: Let $X_k^{(n)} = \mu^{(n)} + U_k^{(n)}$, where μ is also a Gaussian random variable with zero mean and variance ρ/N. $U_k^{(n)}$ is IID Gaussian, with zero mean and variance $(1-\rho)/N$. We analyze the latter model, where entries are conditionally IID given μ. Correlated columns have been analyzed for lasso in this context [16, 17]. The strongest results due to [16] require correlations to decay asymptotically to zero as $1/\log(D)$, while [17] is not strictly comparable since their results are for high-SNR limit. In contrast, we will show that fundamentally, up to constant correlation can be tolerated. The following theorem provides an upper bound to the probability of error for exact support recovery.

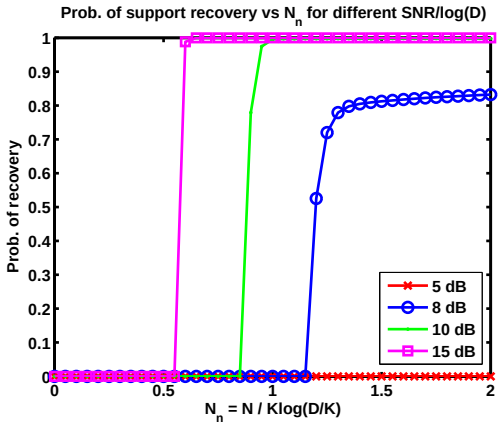

Figure 2: Illustration of SNR cutoff, $K = 32$, $D = 512$.

Theorem 3.2. $P(E) \le \sum_{i=1}^{K} 2^{-Nf(\rho)}$, *where* $f(\rho) = \frac{1}{2}\log\left(1 + (1-\rho)\frac{2i\sigma^2 SNR}{N}\right) - \frac{i}{4N}\log 4 - \frac{\log\binom{D-K}{i}\binom{K}{i}}{N}$.

The first term in $f(\rho)$ is related to the information between X and Y via SNR and ρ, while the second term is related to the uncertainty of β_S and the last term to the uncertainty of S. Note that the error bound given above precisely characterizes the achievable error for any (N, K, D), in contrast to the setting for the sufficient condition in Section 2.1 where K is fixed w.r.t. D. Using this bound, we show the probability of recovery vs. other interesting quantities (see Figs. 4, 5). Also note the relation between $f(\rho)$ and ρ, e.g., for the degenerate case where $\rho = 1$, $f(\rho)$ is negative for any N. This is expected since recovery is not possible in that case, which we prove with the necessity bound.

We now present necessary and sufficient conditions for exact support recovery. We start with a lemma describing the mutual information for this model.

Lemma 3.1.

$$I(X_{\mathcal{S}^1}; Y | X_{\mathcal{S}^2}, \beta_S, \mu) = \frac{1}{2}E\left[\log\left(1 + (1-\rho)\frac{\|\beta_{\mathcal{S}^1}\|^2 SNR}{N}\right)\right],$$

where the expectation is w.r.t. $\beta_{\mathcal{S}^1}$.

The mutual information formula along with the bound given by Theorem 3.2 allow us to obtain the following necessary & sufficient condition for exact recovery.

The necessary condition on SNR is also illustrated in Figure 2, where we plot the probability of error bound given by Theorem 3.2 for different SNR values. Indeed, we show an SNR cutoff regardless of number of measurements as well as tradeoffs beyond the cutoff point. Note that the relation between SNR and N is not explicitly described for lasso [16,17].

Both upper and lower bounds hold for the general case $K = O(D)$, since we use the error bound in Theorem 3.2 to obtain the upper bound instead of Theorem 2.2.

Remark 3.1. *It follows that our relatively simple analysis gives us a bound asymptotically identical to the best-known bound $N = \Omega(K \log(D/K))$ [21] with an independent Gaussian sensing matrix. Our results also incorporate correlations to explicitly characterize the effect of correlated columns on sample complexity. We have shown that the number of samples increases by $\frac{1}{\log(1+(1-\rho)C)}$ relative to $\frac{1}{\log(1+C)}$ for the independent model for some constant C.*

3.2 Noisy variables

We also analyze the additive noise model considered in [14,26], where in the sparse linear regression model (6), a matrix Z^N is observed instead of the sensing matrix X^N, with the relation $Z^N = X^N + V^N$, where $V^{(n)} \sim \mathcal{N}(0, \frac{\nu}{N} I_D)$ IID for $n = 1, \ldots, N$. The rest of the setup is as given in Section 3.1. Note that in contrast to Section 3.1, the model described here exhibits a nonlinear relationship between the variables Z^N and the observations Y^N.

For this problem with noisy observations of variables, we have the following theorem for an upper bound on the probability of error of exact support recovery.

Theorem 3.3. *The error probability of exact support recovery in the noisy data model is given by $P(E) \leq \sum_{i=1}^{K} 2^{-Nf(\rho,\nu)}$, where $f(\rho,\nu) = \frac{1}{2}\log\left(1 + \frac{1-\rho}{1+\nu}\frac{2i\sigma^2 SNR}{N\xi}\right) - \frac{i}{4N}\log 4 - \frac{\log\binom{D-K}{i}\binom{K}{i}}{N}$, where $\xi = 1 + \frac{(1-\rho)\nu}{1+\nu}\frac{KSNR\sigma^2}{N}$.*

The error exponent $f(\rho,\nu)$ differs from $f(\rho)$ defined in Section 3.1 mainly by an extra $1+\nu$ term in the denominator in the log term and reduces to $f(\rho)$ for $\nu = 0$. Also note that $\xi \approx 1$ for sufficiently large N.

We now state a sufficient condition on the number of measurements with the theorem below, which follows from an analysis of the upper bound on recovery error provided in Theorem 3.3.

Theorem 3.4. *For $SNR = \Omega(\log D)$, a sufficient condition on the number of measurements is $N = \Omega\left(\frac{K\log(D/K)}{\log\left(1+\frac{1-\rho}{1+\nu}\sigma^2\right)}\right)$.*

Remark 3.2. *We observe that the sufficient number of measurements is affected by a factor of $\frac{1}{\log(1+C/(1+\nu))}$ in our results, which greatly improves upon the bound with a factor of $(1+\nu)^2$ by [26].*

We also note that while [16,19] analyze correlated Gaussians and others noisy or missing data [14,26] separately, our error bounds and asymptotic sample complexity results unify these into a single expression.

4 Experiments and Comparison with Achievable Bounds

In this section we compare the bounds on the probability of successful recovery we derived in the above sections with the frequency of successful exact support recovery for two recovery algorithms.

For all experiments and evaluation of bounds, we set $K = 32$ and $D = 512$. The variables X^N and observations Y^N are generated according to the normalized model given by (6), where we choose S uniformly at random and let $\beta_S \in \{-1, 1\}^K$ with uniform probability. $N_n = N/(K\log(D/K))$ is the normalized number of measurements.

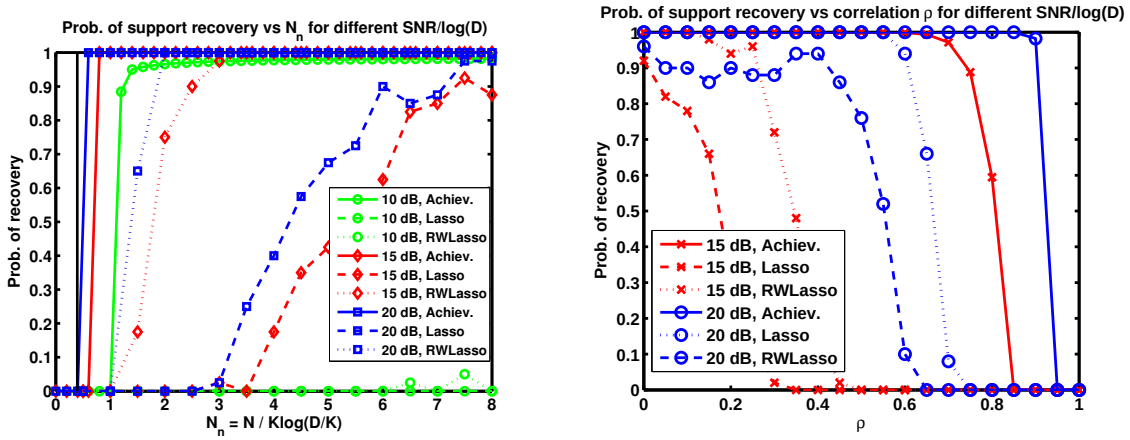

(a) Measurement Cutoffs

(b) Correlation Cutoffs

Figure 3: Comparison of information-theoretic bound vs. lasso and reweighted lasso.

4.1 Lasso and iteratively reweighted lasso

We compare our bounds for independent and correlated sensing elements with lasso [16,17], as defined in [16]. Formally, lasso gives the solution to the following optimization problem:

$$\beta^\star = \arg\min_\beta \frac{1}{2}\left\|Y^N - X^N\beta\right\|_2^2 + \lambda\|\beta\|_1.$$

We set the regularization parameter as $\lambda = 2\sqrt{2\log D}/\sqrt{\mathrm{SNR}}$ as suggested in [16]. We also investigated different values however we have not observed any significant improvements in performance.

We also investigate the performance of a non-convex iterative lasso variant called the iteratively reweighted lasso. This method is proposed in [32] for the noiseless recovery problem; we use an extension for the noisy case, which iteratively solves the following optimization problem at each step:

$$\beta^{(l)} = \arg\min_\beta \frac{1}{2}\left\|Y^N - X^N\beta\right\|_2^2 + \lambda_r \sum_{n=1}^N w_n^{(l)}|\beta_n|, \quad w_n^{(l)} = \frac{1}{\left|\beta_n^{(l-1)}\right| + \epsilon}.$$

This optimization is the same as lasso except for the individual weights $w_n^{(l)}$ for each component β_n, which depend on the output of the previous iteration. ϵ is a suitably small constant that stabilizes the weights for $|\beta_n|$ close to zero. Setting $\beta^{(0)}$ to the solution of regular lasso, the algorithm iterates until $\|\beta^{(l)} - \beta^{(l-1)}\|$ is smaller than a tolerance constant or a maximum number of iterations is reached.

Reweighted lasso aims to sparsify the estimated β compared to regular lasso. At each iteration, it places greater weight on small variables to sparsify the solution, while the influence of large variables is reduced in order to allow for more sensitivity in identifying the other variables. The authors in [32] intuitively justify the sparsifying properties of the algorithm by noting that iteratively solving the reweighted ℓ_1 problem is a Majorization-Minimization algorithm for the log-sum penalty problem, where the penalty is defined as $\sum_{n=1}^N \log\left(|\beta_n| + \epsilon\right)$. The sparsity encouraging properties of this method can be intuitively justified by the fact that $\log\left(|\beta_n| + \epsilon\right)$ approximates the ℓ_0 penalty much better than ℓ_1 does. It should be noted that the log-sum penalty is non-convex, therefore the iterative reweighted minimization is not guaranteed to converge to its global minimum. Furthermore, [32] notes that small values of ϵ (leading to a better approximation of the ℓ_0 penalty) makes it more likely that the algorithm gets stuck at undesirable local minima.

We have chosen reweighted lasso for comparison with the information-theoretic bound since for CS, the optimal ML decoder can be equivalently written as an ℓ_0 constrained least squares minimization for fixed

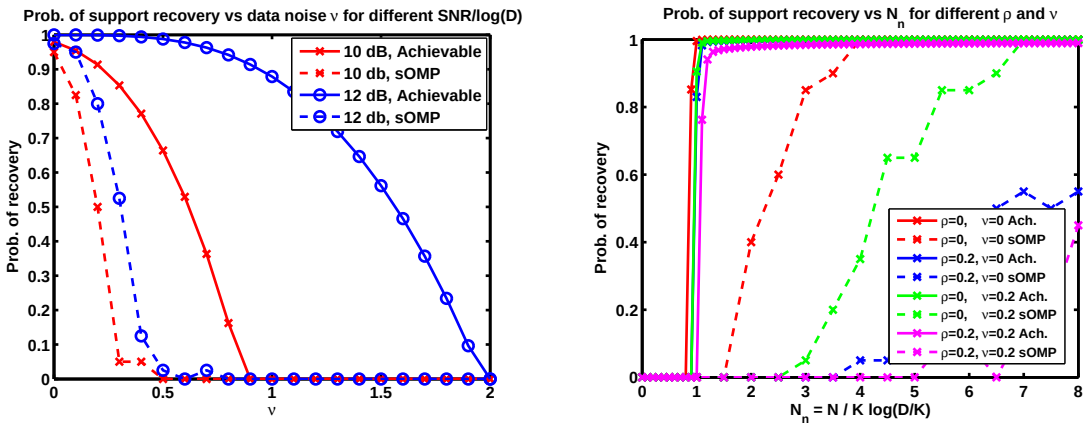

(a) Noise Variance Cutoffs

(b) Effect of Both Noisy Data & Correlation

Figure 4: Comparison of information-theoretic bound vs. support-OMP.

β_S. Therefore we would expect a method like the reweighted lasso to better approach the achievable bound compared to lasso, as it aims to successively approximate the ℓ_0 penalty while still being computationally efficient. We demonstrate that this is the case in our simulation results below.

Figure 3(a) plots the recovery bound for IID variables vs. lasso and reweighted simulation performance, for different number of measurements N. The probabilities of recovery for the lassos are computed over 40 iterations. Compared to lasso, our IT bound has a much sharper transition, while also being tighter, matching closely our lower bound (vertical line for $\mathrm{SNR}/\log D = 20$ dB) obtained with Theorem 3.1. Interestingly, reweighted lasso nearly achieves our performance bounds for high SNR, however it fails in low SNR performance similar to lasso. Note that the theoretical results in [17,19] for lasso are not strictly comparable since they require a significantly large SNR regime. Furthermore, the performance gap approaches infinity as we let K approach D, implying lasso works strictly in sublinear regime.

Figure 3(b) shows our probability of error bound vs. lasso performances for different values of the correlation coefficient ρ, where $N_n = 8$. The probabilities of recovery for lassos are computed over 50 iterations. This plot demonstrates clearly that while our bounds show tolerance to correlation up to a constant approaching 1 (as seen from the sample complexity bound in Theorem 3.1), lasso can tolerate at most $\rho = 0.5$ correlation for exact recovery in this scenario, with very high SNR and N. Note that strongest results due to [16] require correlations to decay asymptotically to zero as $1/\log(D)$. Reweighted lasso shows better performance than lasso, however there is still a significant gap between the achievable correlation bound and the reweighted lasso performance, especially at 15 dB SNR.

4.2 Support-OMP and noisy variables

For the second set of experiments, we compare with a variant of the orthogonal matching pursuit (OMP) algorithm called support-OMP. This algorithm is proposed by [26] and shown to have good performance with theoretical guarantees for problems with noisy or missing observations of the sensing matrix X^N, as we consider in Section 3.2.

Figure 4(a) shows the performance of support-OMP vs. information-theoretic bound, for noisy variables with different noise variances ν. For support-OMP, the recovery probability is computed over 40 iterations. It can be seen that support-OMP performs reasonably well for noisy variables but fails in high variance noise, whereas our information-theoretic bound shows that recovery in much higher noise levels are achievable, especially with higher SNR.

A similar conclusion can be reached from Figure 4(b), where we plot recovery performance for both

correlated and noisy variables. For support-OMP, the recovery probability is computed over 25 iterations. The gap is more pronounced for correlated variables compared to noisy variables, which shows support-OMP is highly affected by correlation and by variable noise to a lesser degree.

5 Discussion

We have presented a framework for analyzing sparse recovery problems that unifies linear and nonlinear observation models, dependent and non-Gaussian measurements matrices, and noisy data. This framework leads to a tight, exponential upper bound on the support recovery error probability and an explicit universal mutual information formula for computing the sample complexity of sparse recovery problems. The central theme here is "inference of a discrete object (sparse support pattern) in a continuous world of observations." We unify sparse problems from an inference perspective by introducing a Markovian assumption. Our approach is not algorithmic and therefore must be used in conjunction with tractable algorithms. Nevertheless, it is useful for identifying gaps between existing algorithms and fundamental information.

Although we consider sparse linear regression and its variants as applications in this paper, there are many other sparse recovery applications for which the framework we consider is applicable and the error bound or the sample complexity bounds we have described are explicitly computable through the formulas in Theorems 2.1 and 2.2. Some examples we have not included due to space considerations are group testing, quantized compressive sensing, multiple regression models or models with missing observations.

As we have shown our approach is also useful in understanding fundamental tradeoffs between different design parameters such as SNR, correlations, measurements matrices and noisy features. For instance, in the linear Gaussian setting we have shown that we could information theoretically tolerate up to constant correlation across different variables while existing results require vanishing correlation. The linear setting has also identified large sample complexity gaps between lasso, support-OMP and information theoretic bounds. Specifically, these gaps get larger as correlation and variable noise increases.

Appendix

We use lower-case $p(Y|X_S)$ notation for the conditional outcome distribution given the true subset of variables averaged over the latent variable β_S. In some cases when we would like to distinguish between the outcome distribution conditioned on different sets of variables we use $p_\omega(\,\cdot\,|\,\cdot\,)$ notation, to emphasize that the conditional distribution is conditioned on the given variables, assuming the true set S is S_ω. W.l.o.g. we assume the true set is S_1 for below proofs. Define $\mathcal{I} = \{1, \ldots, \binom{D}{K}\}$ as the collection of sets ω of size K.

Proof of Theorem 2.1

First, note that $P(E) \leq \sum_{i=1}^{K} P(E_i)$, for E and E_i as defined. If we show separately for each i and any $0 \leq \delta \leq 1$ that the following bound holds, then the theorem follows:

$$P(E_i) \leq 2^{-N\left(E_o(\delta) - \delta \frac{\log \binom{D-K}{i}\binom{K}{i}}{N}\right)}. \tag{A.1}$$

Instead of the above bound, we prove a slightly weaker bound for expositional clarity, which is

$$P(E_i) \leq 2^{-N\left(E_o(\delta) - \frac{\log \binom{D-K}{i}\binom{K}{i}}{N}\right)}. \tag{A.2}$$

Note that the main difference between the above equation and the previous bound is the missing δ term multiplying the binomial expression. The main result follows along the same lines and we refer the reader to [9] for further details.

To prove this result we denote by $\mathcal{A}_i$ the set of indices corresponding to sets of K variables that differ from the true set S_1 in exactly i variables, i.e.,

$$\mathcal{A}_i = \{\omega \in \mathcal{I} : |S_{1^c,\omega}| = i, |S_\omega| = K\} \tag{A.3}$$

We can establish that,

$$\begin{aligned}
\Pr[E_i|\omega_0 = 1, X_{S_1}^N, Y^N, \theta] &\leq \sum_{\omega \in \mathcal{A}_i} \sum_{X_{S_{1^c,\omega}}^N} P(X_{S_{1^c,\omega}}^N|\theta) \frac{p_\omega(Y^N|X_{S_{1,\omega}}^N, X_{S_{1^c,\omega}}^N)^s}{p_1(Y^N|X_{S_{1,\omega}}^N, X_{S_{1,\omega^c}}^N)^s} \qquad \text{(A.4)} \\
&= \sum_{S_{1,\omega}} \sum_{S_{1^c,\omega}} \sum_{X_{S_{1^c,\omega}}^N} P(X_{S_{1^c,\omega}}^N|\theta) \frac{p_\omega(Y^N|X_{S_{1,\omega}}^N, X_{S_{1^c,\omega}}^N)^s}{p_1(Y^N|X_{S_{1,\omega}}^N, X_{S_{1,\omega^c}}^N)^s}.
\end{aligned}$$

Inequality (A.4) is established separately in the following section. It follows that,

$$\begin{aligned}
\Pr[E_i|\omega_0 = 1, X_{S_1}^N, Y^N, \theta] &\leq \left(\sum_{S_{1,\omega}} \sum_{S_{1^c,\omega}} \sum_{X_{S_{1^c,\omega}}^N} P(X_{S_{1^c,\omega}}^N|\theta) \frac{p_\omega(Y^N|X_{S_{1,\omega}}^N, X_{S_{1^c,\omega}}^N)^s}{p_1(Y^N|X_{S_{1,\omega}}^N, X_{S_{1,\omega^c}}^N)^s} \right)^\delta \qquad \text{(A.5)} \\
&\leq \left(\sum_{S_{1,\omega}} \binom{D-K}{i} \sum_{X_{S_{1^c,\omega}}^N} P(X_{S_{1^c,\omega}}^N|\theta) \frac{p_\omega(Y^N|X_{S_{1,\omega}}^N, X_{S_{1^c,\omega}}^N)^s}{p_1(Y^N|X_{S_{1,\omega}}^N, X_{S_{1,\omega^c}}^N)^s} \right)^\delta \qquad \text{(A.6)} \\
&\leq \binom{D-K}{i} \sum_{S_{1,\omega}} \left(\sum_{X_{S_{1^c,\omega}}^N} P(X_{S_{1^c,\omega}}^N|\theta) \frac{p_\omega(Y^N|X_{S_{1,\omega}}^N, X_{S_{1^c,\omega}}^N)^s}{p_1(Y^N|X_{S_{1,\omega}}^N, X_{S_{1,\omega^c}}^N)^s} \right)^\delta, \ \forall s > 0, \ 0 \leq \delta \leq 1. \qquad \text{(A.7)}
\end{aligned}$$

Inequality (A.5) follows from the fact that $\Pr[E_i|\omega_0 = 1, X_{S_1}^N, Y^N, \theta] \leq 1$. Consequently, if U is an upper bound of this probability then it follows that, $\Pr[E_i|\omega_0 = 1, X_{S_1}^N, Y^N, \theta] \leq U^\delta$ for $\delta \in [0, 1]$. Inequality (A.6) follows from symmetry, namely, the inner summation is only dependent on the values of $X_{S_{1^c,\omega}}^N$ and not on the items in the set $S_{1^c,\omega}$. There are exactly $\binom{D-K}{i}$ possible sets $S_{1^c,\omega}$ hence the binomial expression. Note that the sum over $S_{1,\omega}$ cannot be further simplified. This is due to the fact that $X_{S_{1,\omega}}^N$ is already specified since we have conditioned on $X_{S_1}^N$. Since $X_{S_1}^N$ is fixed, the inner sum need not be equal for all sets $S_{1,\omega}, \omega \in \mathcal{A}_i$. Finally, (A.7) follows from standard observation that sum of positive numbers raised to δ-th power for $\delta < 1$ is smaller than the sum of the δ-th power of each number.

We now substitute for the conditional error probability derived above and follow the steps below:

$$
\begin{aligned}
P(E_i) &= \int \sum_{X_{S_1}^N} \sum_{Y^N} P(\theta) P(X_{S_1}^N|\theta) p_1(Y^N|X_{S_1}^N) \Pr[E_i|\omega_0 = 1, X_{S_1}^N, Y^N, \theta] \, \mathrm{d}\theta \\
&\leq \binom{D-K}{i} \int \sum_{S_{1,\omega}} \sum_{Y^N} \sum_{X_{S_1}^N} P(\theta) P(X_{S_1}^N|\theta) p_1(Y^N|X_{S_1}^N) \left(\sum_{X_{S_{1^c,\omega}}^N} P(X_{S_{1^c,\omega}}^N|\theta) \frac{p_\omega(Y^N|X_{S_{1,\omega}}^N, X_{S_{1^c,\omega}}^N)^s}{p_1(Y^N|X_{S_{1,\omega}}^N, X_{S_{1,\omega^c}}^N)^s} \right)^\delta \mathrm{d}\theta
\end{aligned}
$$

Due to symmetry the summation over sets $S_{1,\omega}$ does not depend on ω. Since there are $\binom{K}{K-i}$ sets $S_{1,\omega}$ we get,

$$
\begin{aligned}
P(E_i) &\leq \binom{D-K}{i} \binom{K}{i} \int \sum_{Y^N} \sum_{X_{S_1}^N} P(\theta) P(X_{S_1}^N|\theta) p_1(Y^N|X_{S_1}^N) \left(\sum_{X_{S_{1^c,\omega}}^N} P(X_{S_{1^c,\omega}}^N|\theta) \frac{p_\omega(Y^N|X_{S_{1,\omega}}^N, X_{S_{1^c,\omega}}^N)^s}{p_1(Y^N|X_{S_{1,\omega}}^N, X_{S_{1,\omega^c}}^N)^s} \right)^\delta \mathrm{d}\theta \\
&\leq \binom{D-K}{i} \binom{K}{i} \int \sum_{Y^N} \sum_{X_{S_{1,\omega^c}}^N} \sum_{X_{S_{1,\omega}}^N} P(\theta) P(X_{S_1}^N|\theta) p_1^{1-s\delta}(Y^N|X_{S_{1,\omega}}^N, X_{S_{1,\omega^c}}^N) \\
&\qquad \left(\sum_{X_{S_{1^c,\omega}}^N} P(X_{S_{1^c,\omega}}^N|\theta) p_\omega(Y^N|X_{S_{1,\omega}}^N, X_{S_{1^c,\omega}}^N)^s \right)^\delta \mathrm{d}\theta \\
&= \binom{D-K}{i} \binom{K}{i} \int \sum_{Y^N} \sum_{X_{S_{1,\omega}}^N} P(\theta) P(X_{S_{1,\omega}}^N|\theta) \left(\sum_{X_{S_{1,\omega^c}}^N} P(X_{S_{1,\omega^c}}^N|\theta) p_1^{1/(1+\delta)}(Y^N|X_{S_{1,\omega}}^N, X_{S_{1,\omega^c}}^N) \right)^{1+\delta} \mathrm{d}\theta
\end{aligned}
$$

where the last step follows by letting $s = \frac{1}{1+\delta}$ and noting that from symmetry $X_{S_{1^c,\omega}}^N$ is just a dummy variable and can be replaced by $X_{S_{1,\omega^c}}^N$. This establishes the weaker bound in (A.2), by letting $\mathcal{S}^1 = S_{1,\omega^c}$ and $\mathcal{S}^2 = S_{1,\omega}$.

□

Proof of Equation A.4

Let ζ_ω, $\omega \in \mathcal{A}_i$ denote the event where ω is more likely than 1. Then, from the definition of $\mathcal{A}_i$, the 2 encoded messages differ in i variables. Hence

$$
\Pr[E_i|\omega_0 = 1, X_{S_1}^N, Y^N, \theta] \leq P\Big(\bigcup_{\omega \in \mathcal{A}_i} \zeta_\omega \Big) \leq \sum_{\omega \in \mathcal{A}_i} P(\zeta_\omega)
$$

Now note that $X_{S_1}^N$ shares $(K - i)$ variables with $X_{S_\omega}^N$. Following the introduced notation, the common partition is denoted $X_{S_{1,\omega}}^N$, which is a $N \times (K - i)$ submatrix. The remaining i rows which are in $X_{S_1}^N$

but not in $X^N_{S_\omega}$ are $X^N_{S_{1,\omega^c}}$. Similarly, $X^N_{S_{1^c,\omega}}$ corresponds to variables in $X^N_{S_\omega}$ but not in $X^N_{S_1}$. In other words $X^N_{S_1} = (X^N_{S_{1,\omega}}, X^N_{S_{1,\omega^c}})$ and $X^N_{S_\omega} = (X^N_{S_{1,\omega}}, X^N_{S_{1^c,\omega}})$, where the notation $(F^{N\times n_1}; G^{N\times n_2})$ denotes an $N \times (n_1 + n_2)$ matrix with a submatrix F in the first n_1 columns and G in the remaining n_2 columns. Thus,

$$\begin{aligned} P(\zeta_\omega) &= \sum_{X^N_{S_\omega}: p(Y^N|X^N_{S_\omega}) \geq p(Y^N|X^N_{S_1})} P(X^N_{S_\omega}|X^N_{S_1}, \theta) \\ &\leq \sum_{X^N_{S_{1^c,\omega}}} P(X^N_{S_{1^c,\omega}}|\theta) \frac{p(Y^N|X^N_{S_\omega})^s}{p(Y^N|X^N_{S_1})^s} \quad \forall s > 0,\ \forall \omega \in \mathcal{A}_i \end{aligned} \tag{A.8}$$

where by exchangeability, we have $P(X^N_{S_\omega}|X^N_{\mathcal{S}^1}, \theta) = P(X^N_{S_{1^c,\omega}}|X^N_{\mathcal{S}^1}, \theta) = P(X^N_{S_{1^c,\omega}}|\theta)$ and $\frac{p(Y^N|X^N_{S_\omega})^s}{p(Y^N|X^N_{S_1})^s} \geq 1$ for all $s > 0$, since $\frac{p(Y^N|X^N_{S_\omega})}{p(Y^N|X^N_{S_1})} \geq 1$.

□

Proof of Theorem 2.2

We first derive the sufficiency bound, using the results of Theorem 2.1. To achieve that, we derive a sufficient condition for the error exponent of the error probability $P(E_i)$ in (A.1) to be positive and to drive the error probability to zero as $D \to \infty$. Specifically,

$$Nf(\delta) = NE_o(\delta) - \delta \log \binom{D-K}{i}\binom{K}{i} \to \infty \tag{A.9}$$

where

$$f(\delta) = E_o(\delta) - \delta \frac{\log \binom{D-K}{i}\binom{K}{i}}{N}.$$

To establish the sufficiency bound we follow the argument in [31]. Note that $f(0) = 0$. Since the function $f(\delta)$ is differentiable and has a power series expansion, for a sufficiently small δ, we get by Taylor series expansion in the neighborhood $\delta = 0$ that,

$$f(\delta) = f(0) + \delta \frac{df}{d\delta}\Big|_{\delta=0} + O(\delta^2)$$

Note that

$$\frac{\partial E_o}{\partial \delta}\Big|_{\delta=0} = \frac{I(X^N_{\mathcal{S}^1}; Y^N|X^N_{\mathcal{S}^2}, \theta)}{N}, \tag{A.10}$$

which is shown in the next section.

We can further decompose $I(X^N_{\mathcal{S}^1}; Y^N|X^N_{\mathcal{S}^2}, \theta)$ using the following chain of equalities:

$$\begin{aligned} I(X^N_{\mathcal{S}^1}; Y^N|X^N_{\mathcal{S}^2}, \theta) + I(\beta_S; X^N_{\mathcal{S}^1}|X^N_{\mathcal{S}^2}, Y^N, \theta) &= I(X^N_{\mathcal{S}^1}; Y^N, \beta_S|X^N_{\mathcal{S}^2}, \theta) = I(X^N_{\mathcal{S}^1}; \beta_S|\theta) + I(X^N_{\mathcal{S}^1}; Y^N|X^N_{\mathcal{S}^2}, \beta_S, \theta) \\ &= NI(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2}, \beta_S, \theta), \end{aligned}$$

where the last equality is due to X and β_S being independent and (X^N, Y^N) pairs being independent over n given β_S. Therefore we have

$$\frac{\partial E_o}{\partial \delta}\Big|_{\delta=0} = \frac{I(X^N_{\mathcal{S}^1}; Y^N|X^N_{\mathcal{S}^2}, \theta)}{N} = I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2}, \beta_S, \theta) - \frac{I(\beta_S; X^N_{\mathcal{S}^1}|X^N_{\mathcal{S}^2}, Y^N, \theta)}{N}. \tag{A.11}$$

Now assume that N satisfies

$$N > (1+\epsilon) \frac{\log \binom{D-K}{i}\binom{K}{i}}{I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2}, \beta_S, \theta)}. \tag{A.12}$$

for any constant $\epsilon > 0$. We note that from the Lagrange form of the Taylor Series expansion (an application of the mean value theorem) we can write $E_o(\delta)$ in terms of its first derivative evaluated at zero and a remainder term, i.e.,

$$E_o(\delta) = E_o(0) + \delta E_o'(0) + \frac{\delta^2}{2} E_o''(\psi)$$

for some $\psi \in [0, \delta]$. Hence, for the choice of N in (A.12) and using (A.11) we have

$$Nf(\delta) \geq N \left(\delta \frac{\epsilon}{1+\epsilon} I(X_{\mathcal{S}^1}; Y | X_{\mathcal{S}^2}, \beta_S, \theta) - \delta^2 C I(X_{\mathcal{S}^1}; Y | X_{\mathcal{S}^2}, \beta_S, \theta) - \delta \frac{I(\beta_S; X_{\mathcal{S}^1}^N | X_{\mathcal{S}^2}^N, Y^N, \theta)}{N} \right) \quad \text{(A.13)}$$

where $C = \frac{|E_o''(\psi)|}{2I(X_{\mathcal{S}^1}; Y | X_{\mathcal{S}^2}, \beta_S, \theta)}$ which might depend on K.

A preliminary analysis of the necessary condition that we establish in the next section reveals that $N = \Omega(K \log D)$ is necessary, since $\log \binom{D-K+i}{i} = \Theta(i \log D)$ and $I(X_{\mathcal{S}^1}; Y | X_{\mathcal{S}^2}, \beta_S, \theta) \leq H(Y) = O(1)$. Also, $I(\beta_S; X_{\mathcal{S}^1}^N | X_{\mathcal{S}^2}^N, Y^N, \theta) \leq H(\beta_S)$, which is constant with respect to D since the observation model is only dependent on K variables, due to the sparsity assumption of the observation model $P(Y|X)$. So we see that

$$\frac{I(\beta_S; X_{\mathcal{S}^1}^N | X_{\mathcal{S}^2}^N, Y^N, \theta)}{N} = O\left(\frac{1}{\log D} \right)$$

which is always dominated by $I(X_{\mathcal{S}^1}; Y | X_{\mathcal{S}^2}, \beta_S, \theta)$, which we assumed to be $\omega(1/\log D)$. Therefore we can rewrite (A.13) as

$$Nf(\delta) \geq N \left(\delta \left(\frac{\epsilon}{1+\epsilon} - o(1) \right) I(X_{\mathcal{S}^1}; Y | X_{\mathcal{S}^2}, \beta_S, \theta) - \delta^2 C I(X_{\mathcal{S}^1}; Y | X_{\mathcal{S}^2}, \beta_S, \theta) \right).$$

Finally, if we choose $\delta \leq \frac{\epsilon'}{C}$, where $\epsilon' = \frac{\epsilon}{1+\epsilon}$, then $f(\delta) = \eta$ for some $\eta > 0$ which does not depend on D or N. It follows that $Nf(\delta) \to \infty$ as $D \to \infty$.

We have just shown that for fixed K,

$$N > (1+\epsilon) \cdot \frac{\log \binom{D-K}{i} \binom{K}{i}}{I(X_{\mathcal{S}^1}; Y | X_{\mathcal{S}^2}, \beta_S, \theta)}$$

is sufficient to ensure an arbitrarily small $P(E_i)$. Now note that

$$(1+\epsilon) \binom{D-K+i}{i} \geq \binom{D-K}{i} \binom{K}{i}$$

asymptotically as $D \to \infty$ and K is fixed, for any constant $\epsilon > 0$, which can be incorporated into the previous ϵ as both are arbitrary. Since the average error probability $P(E) \leq \sum_{i=1}^{K} P(E_i)$, it follows that if

$$N > (1+\epsilon) \max_{i=1,\ldots,K} \frac{\log \binom{D-K+i}{i}}{I(X_{\mathcal{S}^1}; Y | X_{\mathcal{S}^2}, \beta_S, \theta)}$$

then for any fixed K, $\lim_{D\to\infty} P(E) = 0$. Consequently, since this is true for any K, $\lim_{K\to\infty} \lim_{D\to\infty} P(E) = 0$.

□

It is important to highlight the main difference between the analysis of the error probability for the problem considered herein and the channel coding problem. In contrast to channel coding, the codewords of a candidate set and the true set are not independent since the two sets could be overlapping. To overcome this difficulty, we separate the error events E_i, $i = 1, \ldots, K$, of misclassifying the true set in i items. Then, for every i we average over realizations of ensemble of codewords for every candidate set while holding fixed the partition common to these sets and the true set of variables.

Proof of Equation A.10

We have

$$E_o(\delta) = -\frac{1}{N}\log \sum_{\theta}\sum_{Y^N}\sum_{X^N_{\mathcal{S}^2}} P(X^N_{\mathcal{S}^2}|\theta)P(\theta)\left[\sum_{X^N_{\mathcal{S}^1}} P(X^N_{\mathcal{S}^1}|\theta)p(Y^N|X^N_{\mathcal{S}^1}, X^N_{\mathcal{S}^2})^{\frac{1}{1+\delta}}\right]^{1+\delta}, \quad 0 \le \delta \le 1$$

where its derivative at $\delta = 0$ can be written as

$$\frac{\partial E_o}{\partial \delta}\Big|_{\delta=0} = -\frac{1}{N}\frac{1}{\sum_{Y^N, X^N_{\mathcal{S}^1}, X^N_{\mathcal{S}^2}, \theta} P(X_{\mathcal{S}^1}|\theta)P(X_{\mathcal{S}^2}|\theta)P(\theta)p(Y^N|X^N_S)}$$
$$\sum_{Y^N, X^N_{\mathcal{S}^2}, \theta} P(X^N_{\mathcal{S}^2}|\theta)P(\theta)\left(\sum_{X^N_{\mathcal{S}^1}} P(X^N_{\mathcal{S}^1}|\theta)p(Y^N|X^N_S)\left[\log\left(\sum_{X^N_{\mathcal{S}^1}} P(X^N_{\mathcal{S}^1}|\theta)p(Y^N|X^N_S)\right) - \log p(Y^N|X^N_S)\right]\right).$$

Noting that $P(X^N_{\mathcal{S}^1}|\theta)P(X^N_{\mathcal{S}^2}|\theta) = P(X^N_{\mathcal{S}^1}, X^N_{\mathcal{S}^2}|\theta)$ by the representation theorem and $p(Y^N|X^N_S) = p(Y^N|X^N_S, \theta)$ by the independence of Y and θ given X_S, above equality simplifies to

$$\begin{aligned}
\frac{\partial E_o}{\partial \delta}\Big|_{\delta=0} &= -\frac{1}{N}\frac{1}{\sum_{Y^N, X^N_{\mathcal{S}^1}, X^N_{\mathcal{S}^2}, \theta} P(Y^N, X^N_S, \theta)} \sum_{Y^N, X^N_{\mathcal{S}^1}, X^N_{\mathcal{S}^2}, \theta} P(Y^N, X^N_S, \theta)\left(\log P(Y^N|X^N_{\mathcal{S}^2}, \theta) - \log P(Y^N|X^N_S, \theta)\right) \\
&= \frac{1}{N}\sum_{Y^N, X^N_{\mathcal{S}^1}, X^N_{\mathcal{S}^2}, \theta} P(Y^N, X^N_S, \theta)\log\frac{P(Y^N, X^N_{\mathcal{S}^2}|X^N_{\mathcal{S}^1}, \theta)}{P(Y^N, X^N_{\mathcal{S}^2}|\theta)} \\
&= \frac{I(X^N_{\mathcal{S}^1}; X^N_{\mathcal{S}^2}, Y^N|\theta)}{N} \\
&= \frac{I(X^N_{\mathcal{S}^1}; Y^N|X^N_{\mathcal{S}^2}, \theta)}{N},
\end{aligned}$$

where the second equality follows by noting the first denominator is equal to 1 and by adding and subtracting $\log P(X^N_{\mathcal{S}^2}|\theta)$ inside the parenthesis. The third equality follows from the definition of mutual information. The final equality follows from the independence of $X_{\mathcal{S}^1}$ and $X_{\mathcal{S}^2}$ given θ.

Necessity bound

The vector of outcomes Y^N is probabilistically related to the index $\omega \in \mathcal{I} = \{1, 2, \ldots, \binom{D}{K}\}$. Suppose $K - i$ elements of the salient set are revealed to us, denoted by $\mathcal{S}^2$. From X^N and Y^N we estimate the set index ω. Let the estimate be $\hat{\omega} = g(X^N, Y^N)$. Define the probability of error $P_e = P(E) = \Pr[\hat{\omega} \neq \omega]$.

E is a binary random variable that takes the value 1 in case of an error i.e., if $\hat{\omega} \neq \omega$, and 0 otherwise, then using the chain rule of entropies [3] we have

$$\begin{aligned}
H(E, \omega|Y^N, X^N, \mathcal{S}^2) &= H(\omega|Y^N, X^N, \mathcal{S}^2) + H(E|\omega, Y^N, X^N, \mathcal{S}^2) \\
&= H(E|Y^N, X^N, \mathcal{S}^2) + H(\omega|E, Y^N, X^N, \mathcal{S}^2). \qquad \text{(A.14)}
\end{aligned}$$

The random variable E is fully determined given X^N, Y^N, ω and $\mathcal{S}^2$. It follows that $H(E|\omega, Y^N, X^N, \mathcal{S}^2) = 0$. Since E is a binary random variable $H(E|Y^N, X^N, \mathcal{S}^2) \le 1$. Consequently, we can bound $H(\omega|E, Y^N, X^N, \mathcal{S}^2)$ as follows,

$$\begin{aligned}
H(\omega|E, Y^N, X^N, \mathcal{S}^2) &= P(E=0)H(\omega|E=0, Y^N, X^N, \mathcal{S}^2) + P(E=1)H(\omega|E=1, Y^N, X^N, \mathcal{S}^2) \\
&\le (1 - P_e)\,0 + P_e \log\left(\binom{D-K+i}{i} - 1\right) \\
&\le P_e \log\binom{D-K+i}{i}. \qquad \text{(A.15)}
\end{aligned}$$

The first inequality follows from the fact that revealing $K-i$ entries, and given that $E=1$, the conditional entropy can be upper bounded by the logarithm of the number of outcomes. From (A.14), we obtain the genie aided Fano's inequality

$$H(\omega|Y^N, X^N, \mathcal{S}^2) \leq 1 + P_e \log \binom{D-K+i}{i} \tag{A.16}$$

Note that for the left hand term, we have

$$\begin{aligned}
H(\omega|Y^N, X^N, \mathcal{S}^2) &= H(\omega|\mathcal{S}^2) - I(\omega; Y^N, X^N|\mathcal{S}^2) \\
&= H(\omega|\mathcal{S}^2) - I(\omega; X^N|\mathcal{S}^2) - I(\omega; Y^N|X^N, \mathcal{S}^2) \\
&\overset{(a)}{=} H(\omega|\mathcal{S}^2) - I(\omega; Y^N|X^N, \mathcal{S}^2) \\
&\overset{(b)}{=} H(\omega|\mathcal{S}^2) - (H(Y^N|X^N, \mathcal{S}^2) - H(Y^N|X^N, \omega)) \\
&\overset{(c)}{=} H(\omega|\mathcal{S}^2) - (H(Y^N|X^N, \mathcal{S}^2, \theta) - H(Y^N|X^N, \omega, \theta)) \\
&\overset{(d)}{\geq} H(\omega|\mathcal{S}^2) - (H(Y^N|X^N_{\mathcal{S}^2}, \theta) - H(Y^N|X^N_{S_\omega}, \theta)) \\
&\overset{(e)}{=} H(\omega|\mathcal{S}^2) - I(X^N_{\mathcal{S}^1}; Y^N|X^N_{\mathcal{S}^2}, \theta)
\end{aligned}$$

where (a) follows from the fact that X^N is independent of $\mathcal{S}^2$ and ω; (b) follows from the fact that conditioning with respect to ω includes conditioning with respect to $\mathcal{S}^2$; (c) follows from the independence of Y and θ given X; (d) follows from the fact that Y^N depends on $\mathcal{S}^2$ only through $X^N_{\mathcal{S}^2}$ and similarly for the second term Y^N depends on ω only through $X^N_{S_\omega}$; the argument for (e) follows by definition.

From (A.16), it then follows that

$$H(\omega|\mathcal{S}^2) - I(X^N_{\mathcal{S}^1}; Y^N|X^N_{\mathcal{S}^2}, \theta) \leq 1 + P_e \log \binom{D-K+i}{i}$$

and since the set $\mathcal{S}^2$ of $K-i$ variables is revealed, ω is uniformly distributed over the set of indices that correspond to sets of size K containing $\mathcal{S}^2$. It follows that

$$\log \binom{D-K+i}{i} - I(X^N_{\mathcal{S}^1}; Y^N|X^N_{\mathcal{S}^2}, \theta) \leq 1 + P_e \log \binom{D-K+i}{i}.$$

Rewriting the above inequality, we have

$$P_e \geq 1 - \frac{I(X^N_{\mathcal{S}^1}; Y^N|X^N_{\mathcal{S}^2}, \theta) + 1}{\log \binom{D-K+i}{i}}. \tag{A.17}$$

Thus, for the probability of error to be asymptotically bounded away from zero, it is necessary that

$$\log \binom{D-K+i}{i} \leq I(X^N_{\mathcal{S}^1}; Y^N|X^N_{\mathcal{S}^2}, \theta) = NI(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2}, \beta_S, \theta) - I(\beta_S; X^N_{\mathcal{S}^1}|X^N_{\mathcal{S}^2}, Y^N, \theta). \tag{A.18}$$

Using (A.11), we can see that

$$N \geq \max_{i=1,\ldots,K} \frac{\log \binom{D-K+i}{i}}{I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2}, \beta_S, \theta) - \frac{I(\beta_S; X^N_{\mathcal{S}^1}|X^N_{\mathcal{S}^2}, Y^N, \theta)}{N}}$$

is a necessary condition for the number of samples N. Finally, since $I(\beta_S; X^N_{\mathcal{S}^1}|X^N_{\mathcal{S}^2}, Y^N) \geq 0$, the following expression is a lower bound to the expression above, proving that it is a necessary condition for recovery,

$$N \geq \max_{i=1,\ldots,K} \frac{\log \binom{D-K+i}{i}}{I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2}, \beta_S, \theta)}.$$

□

Continuous Variables

Even though the results and proof ideas that were used in the above sections are fairly general, the proofs provided for sufficiency bounds were stated for discrete variables and outcomes. In this section we make the necessary generalizations to extend these proofs to continuous variable and observation models. We follow the methodology in [31] and [33].

To simplify the exposition, we consider the extension to continuous variables in the special case of fixed and known β_S and i.i.d. variables. Let $Q(X) = \prod_{i=1}^{D} Q(X_i)$ denote the joint distribution of variables X. The extensions to random β_S and conditionally i.i.d. variables are straightforward. In this case, $I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2}, \beta_S)$ reduces to $I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2})$ and $E_o(\delta)$ reduces to

$$E_o(\delta) = -\log \sum_{Y} \sum_{X_{\mathcal{S}^2}} \left[\sum_{X_{\mathcal{S}^1}} Q(X_{\mathcal{S}^1}) p(Y, X_{\mathcal{S}^2}|X_{\mathcal{S}^1})^{\frac{1}{1+\delta}} \right]^{1+\delta} \qquad 0 \leq \delta \leq 1 \tag{A.19}$$

with $\frac{\partial E_o(\delta)}{\partial \delta}\Big|_{\delta=0} = I(X_{\mathcal{S}^1}; X_{\mathcal{S}^2}, Y) = I(X_{\mathcal{S}^1}; Y|X_{\mathcal{S}^2})$, since $(X^{(n)}, Y^{(n)})$ pairs are independent across n for fixed β_S.

Assume the continuous joint variable probability density $Q(X)$ with joint cumulative density function F and the conditional probability density $p(Y = y|X_S = x)$ for the observation model, which is assumed to be a continuous function of both x and y.

Let $X' \in \mathcal{X}'^D$ be the random vector and $Y' \in \mathcal{Y}'$ be the random variable generated by the quantization of $X \in \mathcal{X}^D = \mathbb{R}^D$ and $Y \in \mathcal{Y} = \mathbb{R}$ respectively, where each variable in X is quantized to L values and Y quantized to J values. Let F' be the joint cumulative density function of X'. As before, let $\hat{S}(X^N, Y^N)$ be the ML decoder with continuous inputs with probability of making i errors in decoding denoted by $P(E_i)$. Let $\hat{S}(X'^N, Y'^N)$ be the ML decoder that quantizes inputs X^N and Y^N to X'^N and Y'^N and has a corresponding probability of error $P'(E_i)$. Define

$$E_o(\delta, X', Y') = -\log \sum_{y' \in \mathcal{Y}'} \sum_{x'_{\mathcal{S}^2} \in \mathcal{X}'^{K-i}} \left[\sum_{x'_{\mathcal{S}^1} \in \mathcal{X}'^i} Q(x'_{\mathcal{S}^1}) p(y', x'_{\mathcal{S}^2}|x'_{\mathcal{S}^1})^{\frac{1}{1+\delta}} \right]^{1+\delta},$$

$$E_o(\delta, X, Y) = -\log \int_{\mathcal{Y}} \int_{\mathcal{X}^{K-i}} \left[\int_{\mathcal{X}^i} Q(x_{\mathcal{S}^1}) p(y, x_{\mathcal{S}^2}|x_{\mathcal{S}^1})^{\frac{1}{1+\delta}} \, \mathrm{d}x_{\mathcal{S}^1} \right]^{1+\delta} \mathrm{d}x_{\mathcal{S}^2} \, \mathrm{d}y,$$

where the indexing denotes the random variates which the error exponents are computed with respect to.

Utilizing the results in the proof of Theorem 3.1 for the discrete models, we will show the following for the continuous model

$$P(E_i) \leq 2^{-N\left(E_o(\delta,X,Y) - \delta \frac{\log \binom{D-K}{i}\binom{K}{i}}{N}\right)}. \tag{A.20}$$

The rest of the proof will then follow as in the discrete case, by noting that $\frac{\partial E_o(\delta,X,Y)}{\partial \delta}\Big|_{\delta=0} = I(X_{\mathcal{S}^1}; X_{\mathcal{S}^2}, Y)$, with the mutual information definition for continuous variables [3].

Our strategy will be the following: we will increase the number of quantization levels for Y' and X' respectively and since discrete result (A.1) holds for any number of quantization levels, by taking limits we will be able to show that

$$P'(E_i) \leq 2^{-N\left(E_o(\delta,X,Y) - \delta \frac{\log \binom{D-K}{i}\binom{K}{i}}{N}\right)}. \tag{A.21}$$

Since $\hat{S}(X^N, Y^N)$ is the minimum probability of error decoder, any upper bound for $P'(E_i)$ will also be an upper bound for $P(E_i)$, proving (A.20).

Assume Y is quantized with the quantization boundaries denoted by $a_1, \ldots, a_{J-1}$, with $Y' = a_j$ if $a_{j-1} < Y \leq a_j$. For convenience denote $a_0 = -\infty$ and $a_J = \infty$. Furthermore assume quantization

boundaries are equally spaced, i.e. $a_j - a_{j-1} = \Delta_J$ for $2 \le j \le J-1$. Now we can write the following

$$\begin{aligned}
E_o(\delta, X', Y') &= -\log \sum_{j=1}^{J} \sum_{x'_{\mathcal{S}^2}} \left[\sum_{x'_{\mathcal{S}^1}} Q(x'_{\mathcal{S}^1}) \left(\int_{a_{j-1}}^{a_j} p(y, x'_{\mathcal{S}^2} | x'_{\mathcal{S}^1}) \, \mathrm{d}y \right)^{\frac{1}{1+\delta}} \right]^{1+\delta} \\
&= -\log \Bigg\{ \sum_{j=2}^{J-1} \Delta_J \sum_{x'_{\mathcal{S}^2}} \left[\sum_{x'_{\mathcal{S}^1}} Q(x'_{\mathcal{S}^1}) \left(\frac{\int_{a_{j-1}}^{a_j} p(y, x'_{\mathcal{S}^2} | x'_{\mathcal{S}^1}) \, \mathrm{d}y}{\Delta_J} \right)^{\frac{1}{1+\delta}} \right]^{1+\delta} \\
&\quad + \sum_{x'_{\mathcal{S}^2}} \left[\sum_{x'_{\mathcal{S}^1}} Q(x'_{\mathcal{S}^1}) \left(\int_{-\infty}^{a_1} p(y, x'_{\mathcal{S}^2} | x'_{\mathcal{S}^1}) \, \mathrm{d}y \right)^{\frac{1}{1+\delta}} \right]^{1+\delta} \\
&\quad + \sum_{x'_{\mathcal{S}^2}} \left[\sum_{x'_{\mathcal{S}^1}} Q(x'_{\mathcal{S}^1}) \left(\int_{a_{J-1}}^{\infty} p(y, x'_{\mathcal{S}^2} | x'_{\mathcal{S}^1}) \, \mathrm{d}y \right)^{\frac{1}{1+\delta}} \right]^{1+\delta} \Bigg\}.
\end{aligned}$$

Let $J \to \infty$ and for each J choose the sequence of quantization boundaries such that $\lim \Delta_J = 0$, $\lim a_{J-1} = \infty$, $\lim a_1 = -\infty$. Then the last two terms disappear and using the fundamental theorem of calculus, we obtain

$$\lim_{J \to \infty} E_o(\delta, X', Y') = E_o(\delta, X', Y) = -\log \int_{\mathcal{Y}} \sum_{x'_{\mathcal{S}^2}} \left[\sum_{x'_{\mathcal{S}^1}} Q(x'_{\mathcal{S}^1}) p(y, x'_{\mathcal{S}^2} | x'_{\mathcal{S}^1})^{\frac{1}{1+\delta}} \right]^{1+\delta} \mathrm{d}y. \tag{A.22}$$

Although it is not necessary for our proof, it can also be shown that $E_o(\delta, X', Y')$ increases for finer quantizations of Y', therefore $E_o(\delta, X', Y)$ gives the smallest upper bound over $P'(E_i)$ over the quantizations of Y.

We repeat the same procedure for X. Assume each variable X_n in X is quantized with the quantization boundaries denoted by $b_1, \ldots, b_{L-1}$, with $X'_n = b_l$ if $b_{l-1} < X_n \le b_l$. For convenience denote $b_0 = -\infty$ and $b_L = \infty$. Furthermore assume quantization boundaries are equally spaced, i.e. $b_l - b_{l-1} = \Delta_L$ for $2 \le l \le L-1$. Then we can write

$$\begin{aligned}
E_o(\delta, X', Y) &= -\log \int_Y \sum_{l=1}^{L} \left[\sum_{x'_{\mathcal{S}^1}} Q(x'_{\mathcal{S}^1}) \left(\int_{b_{l-1}}^{b_l} p(y, x_{\mathcal{S}^2} | x'_{\mathcal{S}^1}) \, \mathrm{d}x_{\mathcal{S}^2} \right)^{\frac{1}{1+\delta}} \right]^{1+\delta} \mathrm{d}y \\
&= -\log \int_{\mathcal{Y}} \sum_{l=1}^{L} \left[\int_{\mathcal{X}^i} \left(\int_{b_{l-1}}^{b_l} p(y, x_{\mathcal{S}^2} | x_{\mathcal{S}^1}) \, \mathrm{d}x_{\mathcal{S}^2} \right)^{\frac{1}{1+\delta}} \mathrm{d}F'(x_{\mathcal{S}^1}) \right]^{1+\delta} \mathrm{d}y \qquad (A.23) \\
&= -\log \int_{\mathcal{Y}} \Bigg\{ \sum_{l=2}^{L-1} \Delta_L \left[\int_{\mathcal{X}^i} \left(\frac{\int_{b_{l-1}}^{b_l} p(y, x_{\mathcal{S}^2} | x_{\mathcal{S}^1}) \, \mathrm{d}x_{\mathcal{S}^2}}{\Delta_L} \right)^{\frac{1}{1+\delta}} \mathrm{d}F'(x_{\mathcal{S}^1}) \right]^{1+\delta} \\
&\quad + \int_{\mathcal{X}^i} \left(\int_{-\infty}^{b_1} p(y, x_{\mathcal{S}^2} | x_{\mathcal{S}^1}) \, \mathrm{d}x_{\mathcal{S}^2} \right)^{\frac{1}{1+\delta}} \mathrm{d}F'(x_{\mathcal{S}^1}) \\
&\quad + \int_{\mathcal{X}^i} \left(\int_{b_{L-1}}^{\infty} p(y, x_{\mathcal{S}^2} | x_{\mathcal{S}^1}) \, \mathrm{d}x_{\mathcal{S}^2} \right)^{\frac{1}{1+\delta}} \mathrm{d}F'(x_{\mathcal{S}^1}) \Bigg\} \mathrm{d}y.
\end{aligned}$$

where (A.23) follows with $F'(x_{\mathcal{S}^1})$ being the step function which represents the cumulative density function of the quantized variables $X'_{\mathcal{S}^1}$.

Let $L \to \infty$, for each L choose a set of quantization point such that $\lim \Delta_L = 0$, $\lim b_{L-1} = \infty$, $\lim b_1 = -\infty$. Again, the second and third terms disappear and the first sum converges to the integral over $X_{\mathcal{S}^2}$. Note that $p(y, x_{\mathcal{S}^2}|x_{\mathcal{S}^1})$ is a continuous function of all its variables since it was assumed that $Q(x)$ and $p(y|x)$ were continuous. Also note that $\lim_{L\to\infty} F' = F$, which implies the weak convergence of the probability measure of X' to the probability measure of X. Given these facts, using the portmanteau theorem we obtain that $E_{F'}\left[p(Y, X_{\mathcal{S}^2}|X_{\mathcal{S}^1})\right] \to E_F\left[p(Y, X_{\mathcal{S}^2}|X_{\mathcal{S}^1})\right]$, which leads to

$$\lim_{L\to\infty} E_o(\delta, X', Y) = -\log \int_{\mathcal{Y}} \int_{\mathcal{X}^{K-i}} \left[\int_{\mathcal{X}^i} p(y, x_{\mathcal{S}^2}|x_{\mathcal{S}^1})^{\frac{1}{1+\delta}} \, \mathrm{d}F(x_{\mathcal{S}^1})\right]^{1+\delta} \mathrm{d}x_{\mathcal{S}^2} \, \mathrm{d}y = E_o(\delta, X, Y). \quad \text{(A.24)}$$

This leads to the following result, completing the proof.

$$P(E_i) \leq P'(E_i) \leq \lim_{J,L\to\infty} 2^{-N\left(E_o(\delta,X',Y') - \delta \frac{\log \binom{D-K}{i}\binom{K}{i}}{N}\right)} = 2^{-N\left(E_o(\delta,X,Y) - \delta \frac{\log \binom{D-K}{i}\binom{K}{i}}{N}\right)}. \quad \text{(A.25)}$$

Proof of Theorem 3.2

To derive the upper bound on error probability, we compute $E_o(\delta)$ explicitly and replace it in Theorem 2.1. First we compute for the easier case, with fixed $\beta_S = \sigma$. In this case, note that (X, Y) pairs are independent across samples and

$$E_o(\delta) = -\log \int_\theta P(\theta) \int_Y \int_{X_{\mathcal{S}^2}} P(X_{\mathcal{S}^2}|\theta) \left[\int_{X_{\mathcal{S}^1}} P(X_{\mathcal{S}^1}|\theta) p(Y|X_{\mathcal{S}^1}, X_{\mathcal{S}^2})^{\frac{1}{1+\delta}} \, \mathrm{d}X_{\mathcal{S}^1}\right]^{1+\delta} \mathrm{d}X_{\mathcal{S}^2} \, \mathrm{d}Y \, \mathrm{d}\theta, \quad 0 \leq \delta \leq 1.$$

For the correlated Gaussian variables, this reduces to

$$E_o(\delta) = -\log \int_\mu \mathcal{N}(\mu; 0, \rho/N) \int_Y \int_{X_{\mathcal{S}^2}} \mathcal{N}(x_2; (K-i)\mu, (K-i)(1-\rho)/N) \left[\int_{X_{\mathcal{S}^1}} \mathcal{N}(x_1; i\mu, i(1-\rho)/N)\mathcal{N}(y - \sigma(x_1 + x_2); 0, 1/\mathrm{SNR})^{\frac{1}{1+\delta}} \, \mathrm{d}x_1\right]^{1+\delta} \mathrm{d}x_2 \, \mathrm{d}y \, \mathrm{d}\mu.$$

As the first step, we input the Gaussian distributions and take the integral inside the brackets over x_1, which gives us

$$\Big[\int_{X_{\mathcal{S}^1}} \mathcal{N}(x_1; i\mu, i(1-\rho)/N)\mathcal{N}(y - \sigma(x_1 + x_2); 0, 1/\mathrm{SNR})^{\frac{1}{1+\delta}} \, \mathrm{d}x_1\Big]^{1+\delta} = \frac{\left(\sqrt{\frac{1}{\mathrm{SNR}\sigma^2}}\right)^\delta}{\sigma\sqrt{2\pi}\left(\sqrt{\frac{i(1-\rho)}{N(1+\delta)} + \frac{1}{\mathrm{SNR}\sigma^2}}\right)^{1+\delta}} \exp\left(-\frac{(\frac{y}{\sigma} - x_2 - i\mu)^2}{2\left(\frac{i(1-\rho)}{N(1+\delta)} + \frac{1}{\mathrm{SNR}\sigma^2}\right)}\right).$$

By plugging in this expression and integrating over x_2, we then have

$$\int_{X_{\mathcal{S}^2}} \mathcal{N}(x_2; (K-i)\mu, (K-i)(1-\rho)/N)\Big[\int_{X_{\mathcal{S}^1}} \mathcal{N}(x_1; i\mu, i(1-\rho)/N)\mathcal{N}(y - \sigma(x_1 + x_2); 0, 1/\mathrm{SNR})^{\frac{1}{1+\delta}} \, \mathrm{d}x_1\Big]^{1+\delta} \mathrm{d}x_2$$
$$= \frac{1}{\sigma\sqrt{2\pi}} \left(\frac{1}{\sqrt{1 + \frac{i(1-\rho)\mathrm{SNR}\sigma^2}{N(1+\delta)}}}\right)^\delta \frac{1}{\sqrt{\frac{(K-i)(1-\rho)}{N} + \frac{i(1-\rho)}{N(1+\delta)} + \frac{1}{\mathrm{SNR}\sigma^2}}} \exp\left(-\frac{(y - \sigma K\mu)^2}{2\sigma^2\left(\frac{(K-i)(1-\rho)}{N} + \frac{i(1-\rho)}{N(1+\delta)} + \frac{1}{\mathrm{SNR}\sigma^2}\right)}\right)$$

Integrating the above expression over y, we are left with

$$\left(\frac{1}{\sqrt{1 + \frac{i(1-\rho)\mathrm{SNR}\sigma^2}{N(1+\delta)}}}\right)^\delta,$$

which no longer depends on μ, therefore the expectation over μ is equal to the above expression and finally we have

$$E_o(\delta) = \frac{\delta}{2} \log \left(1 + (1-\rho) \frac{i\sigma^2 \mathrm{SNR}}{N(1+\delta)} \right),$$

for any $0 \leq \delta \leq 1$.

Now we will show a lower bound on the error exponent $E_o(\delta)$ for the case where β_S is random and IID $\mathcal{N}(0, \sigma^2)$. In this case, $Y^{(n)}$ are not independent across n. In order to lower bound E_o, we first upper bound the observation probability such that,

$$p(Y^N | X^N_{\mathcal{S}^1}, X^N_{\mathcal{S}^2})^{\frac{1}{1+\delta}} = \left(\int_{\beta_S} P(\beta_S) P(Y^N | X^N_{\mathcal{S}^1}, X^N_{\mathcal{S}^2}, \beta_S) \, \mathrm{d}\beta_S \right)^{\frac{1}{1+\delta}} \leq \int_{\beta_S} P(\beta_S)^{\frac{1}{1+\delta}} P(Y^N | X^N_{\mathcal{S}^1}, X^N_{\mathcal{S}^2}, \beta_S)^{\frac{1}{1+\delta}} \, \mathrm{d}\beta_S$$

which follows from the subadditivity of exponent $\frac{1}{1+\delta}$. A lower bound on E_o is then given by

$$E_o(\delta) \geq -\frac{1}{N} \log M^{1+\delta} \int_{\theta^N} P(\theta^N) \int_{Y^N} \int_{X^N_{\mathcal{S}^2}} P(X^N_{\mathcal{S}^2} | \theta^N) \left[\int_{\beta_S} \frac{P(\beta_S)^{\frac{1}{1+\delta}}}{M} \int_{X^N_{\mathcal{S}^1}} P(X^N_{\mathcal{S}^1} | \theta^N) P(Y^N | X^N_{\mathcal{S}^1}, X^N_{\mathcal{S}^2}, \beta_S)^{\frac{1}{1+\delta}} \, \mathrm{d}X^N_{\mathcal{S}^1} \, \mathrm{d}\beta_S \right]^{1+\delta} \mathrm{d}X^N_{\mathcal{S}^2} \, \mathrm{d}Y^N \, \mathrm{d}\theta$$

where $M = \int P(\beta_S)^{\frac{1}{1+\delta}} \, \mathrm{d}\beta_S$ and then by Jensen's inequality, it follows that

$$E_o(\delta) \geq -\frac{1}{N} \log M^{\delta} \int_{\beta_S} P(\beta_S)^{\frac{1}{1+\delta}} \left(\int_{\theta} P(\theta) \int_{Y} \int_{X_{\mathcal{S}^2}} P(X_{\mathcal{S}^2} | \theta) \left[\int_{X_{\mathcal{S}^1}} P(X_{\mathcal{S}^1} | \theta) P(Y | X_{\mathcal{S}^1}, X_{\mathcal{S}^2}, \beta_S)^{\frac{1}{1+\delta}} \, \mathrm{d}X_{\mathcal{S}^1} \right]^{1+\delta} \mathrm{d}X_{\mathcal{S}^2} \, \mathrm{d}Y \, \mathrm{d}\theta \right)^N \mathrm{d}\beta_S \tag{A.26}$$

where we also used the independence of $(X^{(n)}, Y^{(n)})$ across n given β_S.

We start by taking the integral inside the square brackets. For the linear model set-up we have,

$$\begin{aligned}
\int_{X_{\mathcal{S}^1}} P(X_{\mathcal{S}^1} | \theta) P(Y | X_{\mathcal{S}^1}, X_{\mathcal{S}^2}, \beta_S)^{\frac{1}{1+\delta}} \, \mathrm{d}X_{\mathcal{S}^1} &= \int_{\mathbb{R}^i} \mathcal{N}\left(x; \mu 1_i, \frac{1-\rho}{N} I_i\right) \mathcal{N}\left(y - x^\top \beta_1 - x_2^\top \beta_2; 0, 1/\mathrm{SNR}\right)^{\frac{1}{1+\delta}} \, \mathrm{d}x \\
&= \left(\frac{1}{\sqrt{2\pi A}} \right)^i \left(\frac{\sqrt{\mathrm{SNR}}}{\sqrt{2\pi}} \right)^{\frac{1}{1+\delta}} \int_{\mathbb{R}^i} \exp\left(-\frac{(x - \mu 1_i)^\top (x - \mu 1_i)}{2A} \right) \exp\left(-\frac{(y - x^\top \beta_1 - x_2^\top \beta_2)^2}{2B} \right) \mathrm{d}x \\
&= \left(\frac{1}{\sqrt{2\pi A}} \right)^i \left(\frac{\sqrt{\mathrm{SNR}}}{\sqrt{2\pi}} \right)^{\frac{1}{1+\delta}} \int_{\mathbb{R}^i} \exp\left(-\frac{x^\top x}{2A} - \frac{(x^\top \beta_1 + C)^2}{2B} \right) \mathrm{d}x \\
&= \left(\frac{1}{\sqrt{2\pi A}} \right)^i \left(\frac{\sqrt{\mathrm{SNR}}}{\sqrt{2\pi}} \right)^{\frac{1}{1+\delta}} \int_{\mathbb{R}^i} \exp\left(-\frac{1}{2} (x + (BD)^{-1} AC\beta_1)^\top \frac{D}{A} (x + (BD)^{-1} AC\beta_1) \right) \exp\left(-\frac{C^2}{2E} \right) \mathrm{d}x
\end{aligned}$$

where $A = \frac{1-\rho}{N}$, $B = \frac{1+\delta}{\mathrm{SNR}}$, $C = x_2^\top \beta_2 + \mu 1_i^\top \beta_1 - y$, $D = I_i + \frac{A}{B} \beta_1 \beta_1{}^\top$ and $E = \frac{B}{1 - \frac{A}{B} \beta_1{}^\top D^{-1} \beta_1}$. Then taking the integral, some terms on the left cancel and we have

$$\int_{X_{\mathcal{S}^1}} P(X_{\mathcal{S}^1} | \theta) P(Y | X_{\mathcal{S}^1}, X_{\mathcal{S}^2}, \beta_S)^{\frac{1}{1+\delta}} \, \mathrm{d}X_{\mathcal{S}^1} = \left(\frac{\sqrt{\mathrm{SNR}}}{\sqrt{2\pi}} \right)^{\frac{1}{1+\delta}} \frac{1}{\sqrt{|D|}} \exp\left(-\frac{C^2}{2E} \right). \tag{A.27}$$

Writing the second integral that is over $X_{\mathcal{S}^2} = x_2$, we then have

$$\begin{aligned}
&\int_{X_{\mathcal{S}^2}} P(X_{\mathcal{S}^2}|\theta) \left[\int_{X_{\mathcal{S}^1}} P(X_{\mathcal{S}^1}|\theta) P(Y|X_{\mathcal{S}^1}, X_{\mathcal{S}^2}, \beta_S)^{\frac{1}{1+\delta}} \, \mathrm{d}X_{\mathcal{S}^1} \right]^{1+\delta} \mathrm{d}X_{\mathcal{S}^2} \\
&\quad = \sqrt{\frac{\mathrm{SNR}}{2\pi}} \frac{1}{\sqrt{|D|}^{(1+\delta)}} \int_{\mathbb{R}^{K-i}} \mathcal{N}(x; \mu 1_{K-i}, A I_{K-i}) \exp\left(-\frac{(x^\top \beta_2 + \mu 1_i^\top \beta_1 - y)^2}{2E'} \right) \mathrm{d}x \\
&\quad = \sqrt{\frac{\mathrm{SNR}}{2\pi}} \frac{1}{\sqrt{|D|}^{(1+\delta)}} \left(\frac{1}{\sqrt{2\pi A}} \right)^{K-i} \int_{\mathbb{R}^{K-i}} \exp\left(-\frac{x^\top x}{2A} - \frac{(x^\top \beta_2 + F)^2}{2E'} \right) \mathrm{d}x \\
&\quad = \sqrt{\frac{\mathrm{SNR}}{2\pi}} \frac{1}{\sqrt{|D|}^{(1+\delta)}} \left(\frac{1}{\sqrt{2\pi A}} \right)^{K-i} \int_{\mathbb{R}^{K-i}} \exp\left(-\frac{1}{2}(x + (E'G)^{-1} A F \beta_2)^\top \frac{G}{A} (x + (E'G)^{-1} A F \beta_2) \right) \exp\left(-\frac{F^2}{2H} \right) \mathrm{d}x
\end{aligned}$$

where $E' = \frac{E}{1+\delta}$, $F = \mu 1_K^\top \beta_S - y$, $G = 1 + \frac{A}{E'} \beta_2 {\beta_2}^\top$ and $H = \frac{E'}{1 - \frac{A}{E'} {\beta_2}^\top G^{-1} \beta_2}$. Again, evaluating the integral, we obtain

$$\int_{X_{\mathcal{S}^2}} P(X_{\mathcal{S}^2}|\theta) \left[\int_{X_{\mathcal{S}^1}} P(X_{\mathcal{S}^1}|\theta) P(Y|X_{\mathcal{S}^1}, X_{\mathcal{S}^2}, \beta_S)^{\frac{1}{1+\delta}} \, \mathrm{d}X_{\mathcal{S}^1} \right]^{1+\delta} \mathrm{d}X_{\mathcal{S}^2} = \sqrt{\frac{\mathrm{SNR}}{2\pi}} \frac{1}{\sqrt{|D|}^{(1+\delta)}} \frac{1}{\sqrt{|G|}} \exp\left(-\frac{F^2}{2H} \right).$$

Integrating the above expression w.r.t. $Y = y$, we see that the result is independent of $\theta = \mu$, and therefore

$$\int_\theta P(\theta) \int_Y \int_{X_{\mathcal{S}^2}} P(X_{\mathcal{S}^2}|\theta) \left[\int_{X_{\mathcal{S}^1}} P(X_{\mathcal{S}^1}|\theta) P(Y|X_{\mathcal{S}^1}, X_{\mathcal{S}^2}, \beta_S)^{\frac{1}{1+\delta}} \, \mathrm{d}X_{\mathcal{S}^1} \right]^{1+\delta} \mathrm{d}X_{\mathcal{S}^2} \, \mathrm{d}Y \, \mathrm{d}\theta = \frac{\sqrt{\mathrm{SNR}}}{\sqrt{|D|}^{(1+\delta)}} \sqrt{\frac{H}{|G|}}.$$

By the matrix determinant lemma, we have $|D| = 1 + \frac{A}{B} {\beta_1}^\top \beta_1$ and by the Sherman-Morrison formula, $D^{-1} = I_i - \frac{\beta_1 {\beta_1}^\top}{\frac{B}{A} + {\beta_1}^\top \beta_1}$. Similarly, $|G| = 1 + \frac{A}{E'} {\beta_2}^\top \beta_2$ and $G^{-1} = I_i - \frac{\beta_2 {\beta_2}^\top}{\frac{E'}{A} + {\beta_2}^\top \beta_2}$. By plugging in these expressions, we can then see that $E' = \frac{B|D|}{1+\delta}$ and $H = E'|G|$. We simplify the above expression to obtain

$$\frac{\sqrt{\mathrm{SNR}}}{\sqrt{|D|}^{(1+\delta)}} \sqrt{\frac{H}{|G|}} = \frac{\sqrt{\mathrm{SNR}}}{\sqrt{|D|}^{(1+\delta)}} \sqrt{\frac{B|D|}{1+\delta}} = \left(\frac{1}{\sqrt{|D|}} \right)^\delta = \left(1 + (1-\rho) \frac{\mathrm{SNR} {\beta_1}^\top \beta_1}{N(1+\delta)} \right)^{-\delta/2}. \tag{A.28}$$

Note that this expression is analogous to the bound we obtained for the fixed case, since $E[{\beta_1}^\top \beta_1] = i\sigma^2$. With the above bound, we will now show a lower bound on $E_o(\delta)$ for $\delta = 1$ and $\sigma^2 = \frac{1}{8\pi}$. We note that we choose this σ^2 without loss of generality, since for any value or scaling of σ can be incorporated into the SNR of the problem to obtain an equivalent model, such that $\mathrm{SNR}\sigma^2$ is fixed. This result can also be shown without the assumption on σ^2, but the specific bounding methods we use utilize this assumption. To this effect, we analyze the equivalent problem with parameters $\mathrm{SNR}' = \mathrm{SNR}\sigma^2 8\pi$ and $\sigma'^2 = \frac{1}{8\pi}$. Note that with

this choice of σ'^2 and δ, $M = \int_{\mathbb{R}^K} P(\beta_S)^{\frac{1}{2}} \, \mathrm{d}\beta_S = 1^K = 1$. Using (A.26), we now write,

$$
\begin{aligned}
E_o(1) &\geq -\frac{1}{N} \log M \int_{\mathbb{R}^K} P(\beta_S)^{\frac{1}{2}} \left(1 + (1-\rho)\frac{\mathrm{SNR}{\beta_1}^\top \beta_1}{2N}\right)^{-\frac{N}{2}} \mathrm{d}\beta_S \\
&= -\frac{1}{N} \log \int_{\mathbb{R}^i} P(\beta_1)^{\frac{1}{2}} \left(1 + (1-\rho)\frac{\mathrm{SNR}'{\beta_1}^\top \beta_1}{2N}\right)^{-\frac{N}{2}} \mathrm{d}\beta_1 \\
&= -\frac{1}{N} \log(\sqrt{4})^{\frac{i}{2}} \int_{\mathbb{R}^i} \exp\left[-\frac{{\beta_1}^\top \beta_1}{4\sigma'^2}\right] \left(1 + (1-\rho)\frac{\mathrm{SNR}'{\beta_1}^\top \beta_1}{2N}\right)^{-\frac{N}{2}} \mathrm{d}\beta_1 \\
&\geq -\frac{1}{N} \log(\sqrt{4})^{\frac{i}{2}} \left(\sqrt{8\pi\sigma'^2}\right)^{\frac{iN}{2}} \int_{\mathbb{R}^i} \left[\left(\frac{1}{\sqrt{8\pi\sigma'^2}}\right)^i \exp\left[-\frac{{\beta_1}^\top \beta_1}{8\sigma'^2}\right] \left(1 + (1-\rho)\frac{\mathrm{SNR}'{\beta_1}^\top \beta_1}{2N}\right)\right]^{-\frac{N}{2}} \mathrm{d}\beta_1 \\
&\geq -\frac{1}{N} \log 4^{\frac{i}{4}} \left[\int_{\mathbb{R}^i} \left(\frac{1}{\sqrt{8\pi\sigma'^2}}\right)^i \exp\left[-\frac{{\beta_1}^\top \beta_1}{8\sigma'^2}\right] \left(1 + (1-\rho)\frac{\mathrm{SNR}'{\beta_1}^\top \beta_1}{2N}\right) \mathrm{d}\beta_1\right]^{-\frac{N}{2}} \\
&= \frac{1}{2} \log\left(1 + (1-\rho)\frac{2i\mathrm{SNR}\sigma^2}{N}\right) - \frac{i}{4N} \log 4.
\end{aligned}
$$

The first equality follows by taking β_2 out of the integral and noting that $\int_{\mathbb{R}^{K-i}} P(\beta_2)^{\frac{1}{2}} \, \mathrm{d}\beta_2 = 1^{K-i} = 1$. We obtain the second equality by expanding $P(\beta_1)^{\frac{1}{2}}$. We upper bound $\exp\left[-\frac{{\beta_1}^\top \beta_1}{8\sigma'^2}\right]$ by $\exp\left[-\frac{{\beta_1}^\top \beta_1}{8\sigma'^2}\right]^{-\frac{N}{2}}$ where $0 \leq \exp\left[-\frac{{\beta_1}^\top \beta_1}{8\sigma'^2}\right] \leq 1$ to obtain the first inequality and the second one follows by the superadditivity of exponentiating with $-\frac{N}{2}$. Finally, we note that the integral is an expectation w.r.t. $\beta_1 \sim \mathcal{N}(0, 4\sigma'^2 I_i)$ and obtain the last equality, where we also replace SNR' and σ'^2.

□

Proof of Lemma 3.1

Note the following equalities,

$$
\begin{aligned}
I(X_{\mathcal{S}^1}; Y | X_{\mathcal{S}^2}, \beta_S, \mu) &= h(Y|X_{\mathcal{S}^2}, \beta_S, \mu) - h(Y|X_S, \beta_S, \mu) \\
&= h\left(X_{\mathcal{S}^1}^\top \beta_{\mathcal{S}^1} + W | \beta_{\mathcal{S}^1}, \mu\right) - h(W) \\
&= E_{\beta_{\mathcal{S}^1}, \mu}\left[\frac{1}{2} \ln\left(2\pi e \left(\mathrm{var}\left(X_{\mathcal{S}^1}^\top \beta_{\mathcal{S}^1} | \beta_{\mathcal{S}^1}, \mu\right) \, \beta_{\mathcal{S}^1}^\top \beta_{\mathcal{S}^1} + \frac{1}{\mathrm{SNR}}\right)\right)\right] - \frac{1}{2} \ln\left(2\pi e \, \frac{1}{\mathrm{SNR}}\right) \\
&= E_{\beta_{\mathcal{S}^1}}\left[\frac{1}{2} \ln\left(1 + (1-\rho)\frac{\beta_{\mathcal{S}^1}^\top \beta_{\mathcal{S}^1} \mathrm{SNR}}{N}\right)\right],
\end{aligned}
$$

where the second equality follows from the independence of $X_{\mathcal{S}^1}$ and $X_{\mathcal{S}^2}$ given μ and the last equality follows from the fact that $\mathrm{var}(X_{\mathcal{S}^1}^\top \beta_{\mathcal{S}^1} | \beta_{\mathcal{S}^1}, \mu) = \beta_{\mathcal{S}^1}^\top E[U_{\mathcal{S}^1} U_{\mathcal{S}^1}^\top] \beta_{\mathcal{S}^1} = \beta_{\mathcal{S}^1}^\top \beta_{\mathcal{S}^1} \frac{1-\rho}{N}$.

Proof of Theorem 3.1

We first show that $\mathrm{SNR} = \log D$ is a necessary condition. For any D, K or SNR assume N is much larger such that

$$
E\left[\ln\left(1 + (1-\rho)\frac{\beta_{\mathcal{S}^1}^\top \beta_{\mathcal{S}^1} \mathrm{SNR}}{N}\right)\right] \asymp E\left[(1-\rho)\frac{\beta_{\mathcal{S}^1}^\top \beta_{\mathcal{S}^1} \mathrm{SNR}}{N}\right] = (1-\rho)\frac{i\sigma^2 \mathrm{SNR}}{N}.
$$

Then the necessary condition given by Theorem 3.1 is

$$
N > C \max_i \frac{\log \binom{D-K}{i}\binom{K}{i}}{(1-\rho)\frac{i\sigma^2 \mathrm{SNR}}{N}}
$$

which readily leads to the condition that

$$\mathrm{SNR} > C \max_i \frac{\log \binom{D-K}{i}\binom{K}{i}}{(1-\rho)i\sigma^2} \asymp \log D \tag{A.29}$$

for σ constant.

From the upper bound given by Theorem 3.1, the sufficiency bound in Theorem 3.2 is obtained in a straightforward manner, by looking at conditions where $Nf(\rho)$ goes to infinity. So for each i, we have

$$P(E_i) \leq 2^{-\left(N\frac{1}{2}\log\left(1+(1-\rho)\frac{2i\sigma^2\mathrm{SNR}}{N}\right)-\frac{i}{4}\log 4-\log\binom{D-K}{i}\binom{K}{i}\right)},$$

then, as $\log \binom{D-K}{i}\binom{K}{i} = \Theta(i\log(D/i))$ dominates $\frac{i}{4}\log 4$ we can see that the following is a sufficient condition on N for exact support recovery:

$$N > (1+\epsilon) \max_{i=1,\ldots,K} \frac{2\log \binom{D-K}{i}\binom{K}{i}}{\log\left(1+(1-\rho)\frac{2i\mathrm{SNR}\sigma^2}{N}\right)}. \tag{A.30}$$

Assume $\mathrm{SNR} > C\frac{\log D}{(1-\rho)\sigma^2}$. Also assume $N = \Omega\left(\frac{K\log(D/K)}{\log(1+(1-\rho)\sigma^2)}\right)$, as in the theorem statement. Then, the bound in (A.30) becomes

$$\max_{i=1,\ldots,K} \frac{2\log \binom{D-K}{i}\binom{K}{i}}{\log\left(1+(1-\rho)\frac{2i\mathrm{SNR}\sigma^2}{N}\right)} \asymp \max_{i=1,\ldots,K} \frac{i\log(D/i)}{\log\left(1+2C\frac{i}{K}\log(1+(1-\rho)\sigma^2)\frac{\log D}{\log(D/K)}\right)},$$

where we assume σ^2 constant, w.l.o.g., since the scaling of elements of β_S can instead be incorporated into SNR to obtain an equivalent model as we did in the proof of Theorem 3.1.

First, consider the case $K = o(D)$. Then the sufficient condition reduces to

$$N > \max_{i=1,\ldots,K} \frac{i\log D}{\log\left(1+2C\frac{i}{K}\log(1+(1-\rho)\sigma^2)\right)},$$

which, for the case $i = o(K)$ is

$$N > \frac{i\log D}{C\frac{i}{K}\left(\log(1+(1-\rho)\sigma^2)\right)} \asymp \frac{K\log(D/K)}{\log(1+(1-\rho)\sigma^2)},$$

which is satisfied for chosen N. For $i = \Theta(K)$, asymptotically, we have

$$N > \frac{K\log D}{\log\left(1+2C\log(1+(1-\rho)\sigma^2)\right)} \asymp \frac{K\log(D/K)}{\log\left(1+\log(1+(1-\rho)\sigma^2)\right)},$$

which is also satisfied by N.

Second, consider the case $K = \Theta(D)$. We then have the condition

$$N > \max_{i=1,\ldots,K} \frac{i\log(D/i)}{\log\left(1+2C\frac{i}{K}\log(1+(1-\rho)\sigma^2)\log D\right)},$$

which for $i = o(K)$, is asymptotically equivalent to

$$N > \frac{i\log D}{2C\frac{i}{K}\log(1+(1-\rho)\sigma^2)\log D} = \frac{K}{2C\log(1+(1-\rho)\sigma^2)} \asymp \frac{K\log(D/K)}{\log(1+(1-\rho)\sigma^2)}$$

which is satisfied for chosen N. For $i = \Theta(K)$, asymptotically we have the condition

$$N > \frac{K\log(D/K)}{\log\left(1+\log(1+(1-\rho)\sigma^2)\log D\right)},$$

which is also satisfied for chosen N.

The necessity bound is obtained by using the derived mutual information expression and looking at the case $i = K$. From Lemma 3.1, we have

$$I(X_{\mathcal{S}^1}; Y | X_{\mathcal{S}^2}, \beta_S, \mu) \asymp E_{\beta_{\mathcal{S}^1}} \left[\log \left(1 + (1-\rho) \frac{\beta_{\mathcal{S}^1}^\top \beta_{\mathcal{S}^1} \mathrm{SNR}}{N} \right) \right],$$

which leads to the following necessary condition, as given by Theorem 2.2:

$$N \geq \max_{i=1,\ldots,K} \frac{\log \binom{D-K+i}{i}}{E_{\beta_{\mathcal{S}^1}} \left[\log \left(1 + (1-\rho) \frac{\beta_{\mathcal{S}^1}^\top \beta_{\mathcal{S}^1} \mathrm{SNR}}{N} \right) \right]}.$$

Note that $E_{\beta_{\mathcal{S}^1}} \left[\log \left(1 + (1-\rho) \frac{\beta_{\mathcal{S}^1}^\top \beta_{\mathcal{S}^1} \mathrm{SNR}}{N} \right) \right] \leq \log \left(E_{\beta_{\mathcal{S}^1}} \left[1 + (1-\rho) \frac{\beta_{\mathcal{S}^1}^\top \beta_{\mathcal{S}^1} \mathrm{SNR}}{N} \right] \right) = \log \left(1 + (1-\rho) \frac{i\sigma^2 \mathrm{SNR}}{N} \right)$ due to Jensen's inequality, therefore the following is also a necessary condition, where we consider only $i = K$:

$$N \geq \frac{\log \binom{D}{K}}{\log \left(1 + (1-\rho) \frac{K\sigma^2 \mathrm{SNR}}{N} \right)} \asymp \frac{K \log(D/K)}{\log \left(1 + (1-\rho) \frac{K\sigma^2 \mathrm{SNR}}{N} \right)}. \tag{A.31}$$

Assume $\mathrm{SNR} = \Theta(\log(D/K))$, which is given by (A.29) for $\sigma^2 = O(1)$ and $i = K$. It is then clear that (A.31) does not hold for $N = o(K \log(D/K))$, since $\frac{K \log(D/K)}{N} \geq \log \left(1 + (1-\rho)\sigma^2 \frac{K \log(D/K)}{N} \right)$ for $\sigma^2 = O(1)$. However for $N = \Omega(K \log(D/K))$, the condition (A.31) is

$$N = \Omega \left(\frac{K \log(D/K)}{\log \left(1 + (1-\rho)\sigma^2 \right)} \right),$$

which proves the lower bound in Theorem 3.2.

Proof of Theorem 3.3

To show the upper bound on error probability given in Theorem 3.3, we will write $P(Y | Z_{\mathcal{S}^1}, Z_{\mathcal{S}^2})$ and compute $E_o(\delta)$. For clarity, we consider fixed $\beta_S = \sigma$ as we did initially did in the proof of Theorem 3.1. Note that we can write

$$P(Y | Z_{\mathcal{S}^1}, Z_{\mathcal{S}^2}) = P(Y | Z_{\mathcal{S}^1}, Z_{\mathcal{S}^2}, \mu) = \int_{\mathbb{R}^K} P(Y | X_{\mathcal{S}^1}, X_{\mathcal{S}^2}) P(X_S | Z_S, \mu) \, \mathrm{d}X_S.$$

The first term is given by $\mathcal{N}(y - x^\top \beta_S; 0, 1/\mathrm{SNR})$ as before, for $Y = y$ and $X_S = x$. Let $\alpha = \frac{1}{1+\nu}$, then using the conditional probability of jointly Gaussian random vectors, we have

$$P(X_S = x | Z_S = z, \mu) = \mathcal{N} \left(x; (1-\alpha)\mu 1_K + \alpha z, \frac{1-\rho}{N}(1-\alpha) I_K \right),$$

then, considering only sums of X_S and Z_S as x and z since $\beta_S = \sigma$, as we did in the proof of Theorem 3.1,

the integral is

$$
\begin{aligned}
P\left(Y=y\Big|\sum_{k\in S}Z_k=z,\mu\right) &= \int_{\mathbb{R}}\mathcal{N}(y-\sigma x;0,1/\mathrm{SNR})\,\mathcal{N}\left(x;(1-\alpha)K\mu+\alpha z,\frac{1-\rho}{N}(1-\alpha)K\right)\,\mathrm{d}x \\
&= \frac{1}{2\pi}\sqrt{\frac{1}{AB}}\int_{\mathbb{R}}\exp\left(-\frac{(x-C)^2}{2A}\right)\exp\left(-\frac{(y-\sigma x)^2}{2B}\right)\,\mathrm{d}x \\
&= \frac{1}{2\pi}\sqrt{\frac{1}{AB}}\int_{\mathbb{R}}\exp\left(-\frac{(x-G)^2}{2\frac{AB}{A\sigma^2+B}}\right)\exp\left(-\frac{(y-\sigma C)^2}{2(A\sigma^2+B)}\right)\,\mathrm{d}x \\
&= \frac{1}{2\pi}\sqrt{\frac{1}{AB}}\exp\left(-\frac{(y-\sigma C)^2}{2(A\sigma^2+B)}\right)\sqrt{\frac{AB}{2\pi(A\sigma^2+B)}} \\
&= \sqrt{\frac{1}{2\pi(A\sigma^2+B)}}\exp\left(-\frac{(y-\sigma C)^2}{2(A\sigma^2+B)}\right) \\
&= \mathcal{N}\left(y-\alpha\sigma z-\alpha(1-\alpha)\sigma K\mu;0,\frac{1}{\mathrm{SNR}}+\frac{(1-\rho)(1-\alpha)K\sigma^2}{N}\right),
\end{aligned}
$$

where $A=\frac{1-\rho}{N}(1-\alpha)K$, $B=\frac{1}{\mathrm{SNR}}$, $C=(1-\alpha)\mu K+\alpha z$ and $G=\frac{A\sigma y+BC}{\sqrt{A\sigma^2+B}}$. The last equation follows through the steps used to show (A.27). For the first equality, we compute and replace the probability distributions w.r.t. sums x and z. We obtain the third equality by rewriting the terms inside the exponentials to obtain a square term with x. Then, we take the second exponential outside the integral and compute the integral, which gives us the fourth equality. Note that G does not affect the integration result. Finally in the last step we note that the resulting expression is a Gaussian distribution with the given form.

Note that the resulting probability distribution is the same as $P(Y|X_S)$ we used in the proof of Theorem 3.1 except a few differences: σ is replaced with $\alpha\sigma$, $\frac{1}{\mathrm{SNR}}$ is replaced with $\frac{1}{\mathrm{SNR}}+\frac{(1-\rho)(1-\alpha)K\sigma^2}{N}$ and lastly there is an extra $(1-\alpha)\mu K$ term. This last term does not affect the resulting lower bound on the error exponent E_o, since it disappears in the integration over Y like the other μ terms. We also note that $P(Z_{\mathcal{S}^1}|\mu)$ and $P(Z_{\mathcal{S}^2}|\mu)$ terms in the integral (A.26) are different than $P(X_{\mathcal{S}^1}|\mu)$ and $P(X_{\mathcal{S}^2}|\mu)$. To account for this difference, we need to replace the variance $\frac{1-\rho}{N}$ with $\frac{(1-\rho)(1+\nu)}{N}$ in the integrations w.r.t. z_1 and z_2 that follow.

Finally, by doing the necessary replacements outlined above and following the proof of Theorem 3.1, we obtain the following error exponent for fixed $\beta_S=\sigma$:

$$
E_o(\delta)=\frac{\delta}{2}\log\left(1+\frac{1-\rho}{1+\nu}\frac{i\sigma^2\mathrm{SNR}}{N(1+\delta)\xi}\right),
$$

where $\xi=1+\frac{(1-\rho)\nu}{1+\nu}\frac{K\mathrm{SNR}\sigma^2}{N}$. Then the same analysis in the proof of Theorem 3.1 can be employed for random β_S, to obtain the lower bound,

$$
E_o(1)\geq\frac{1}{2}\log\left(1+\frac{1-\rho}{1+\nu}\frac{2i\sigma^2\mathrm{SNR}}{N\xi}\right)-\frac{i}{4N}\log 4,
$$

which proves the upper bound on error probability given in Theorem 3.3.

Proof of Theorem 3.4

We analyze the upper bound given in Theorem 3.3 to obtain the sufficient condition on N. First, note that $\frac{(1-\rho)\nu}{1+\nu}\leq 1$, therefore $\xi\leq 1+\frac{K\mathrm{SNR}\sigma^2}{N}$.

Let $\mathrm{SNR}=c\log(D/K)$ for now, which is more relaxed than the SNR condition we assume in the theorem and assume $N=\Omega\left(\frac{K\log(D/K)}{\log\left(1+\frac{1-\rho}{1+\nu}\sigma^2\right)}\right)$. Then it is easy to see that $\xi=O(1)$ for $\sigma^2=O(1)$. As before, we can

assume $\sigma^2 = O(1)$ w.l.o.g. since otherwise we can incorporate its scaling into SNR. Then for some constant $C > 0$, we have the lower bound

$$E_o(1) \geq \frac{1}{2}\log\left(1 + \frac{1-\rho}{1+\nu}\frac{2ci\sigma^2\mathrm{SNR}}{CN}\right) - \frac{i}{4N}\log 4,$$

and therefore we have

$$P(E_i) \leq 2^{-\left(N\frac{1}{2}\log\left(1+c'\frac{1-\rho}{1+\nu}\frac{i\sigma^2\mathrm{SNR}}{N}\right) - \frac{i}{4}\log 4 - \log\binom{D-K}{i}\binom{K}{i}\right)},$$

for a constant $c' > 0$.

Following the arguments in the proof of Theorem 3.2, we can see that for N chosen as above, $P(E)$ goes to zero, proving the theorem.

References

[1] D. L. Donoho. Compressed sensing. *IEEE Trans. Inf. Theory*, 52(4):1289 –1306, April 2006.

[2] C. M. Bishop and N. M. Nasrabadi. *Pattern recognition and machine learning*, volume 1. Springer New York, 2006.

[3] T. M. Cover and J. A. Thomas. Elements of information theory, 1991.

[4] M. B. Malyutov and P. S. Mateev. Screening design for non-symmetric response function. *Mat. Zemetki*, 27:109–127, 1980.

[5] M. B. Malyutov. On planning of screening experiments. In *Proceedings of 1975 IEEE-USSR Workshop on Inf. Theory*, 1976.

[6] M. B. Malyutov. The separating property of random matrices. *Mat. Zametki*, 23, 1978.

[7] M. B. Malyutov. Maximal rates of screening designs. *Probability and its Applic.*, 24, 1979.

[8] A. Dyachkov. Lectures on designing screening experiments. *Moscow State Univ.*, 2003.

[9] G. Atia and V. Saligrama. Boolean compressed sensing and noisy group testing. *IEEE Trans. Inf. Theory*, 58(3), March 2012.

[10] C. Aksoylar, G. Atia, and V. Saligrama. Sparse signal processing with linear and non-linear observations: A unified shannon theoretic approach. In *Information Theory Workshop (ITW), 2013 IEEE*, Seville, Spain, 2013.

[11] P. T. Boufounos and R. G. Baraniuk. 1-bit compressive sensing. In *Proc. of Conf. on Information Sciences and Systems (CISS)*, pages 16–21, March 2008.

[12] Y. Plan and R. Vershynin. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. *IEEE Trans. Inf. Theory*, 59(1):482–494, Jan 2013.

[13] S. Negahban and M. Wainwright. Joint support recovery under high-dimensional scaling: Benefits and perils of $\ell_{1,\infty}$-regularization. In *Advances in Neural Information Processing Systems (NIPS)*, 2008.

[14] P.-L. Loh and M. Wainwright. High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity. *arXiv pre-print, arXiv:1109.3714*, September 2011.

[15] C. Aksoylar, G. Atia, and V. Saligrama. Sparse Signal Processing with Linear and Non-Linear Observations: A Unified Shannon Theoretic Approach. *arXiv pre-print, arXiv:1304.0682*, April 2013.

[16] E. J. Candès and Y. Plan. Near-ideal model selection by ℓ_1 minimization. *The Annals of Statistics*, 37(5A):2145–2177, 2009.

[17] M. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1-constrained quadratic programs. In *Allerton Conf. on Communication, Control and Computing*, Monticello, IL, 2006.

[18] C. L. Chan, S. Jaggi, V. Saligrama, and S. Agnihotri. Non-adaptive group testing: Explicit bounds and novel algorithms. In *Proc. of the Int. Symp. on Information Theory (ISIT)*, pages 1837–1841, July 2012.

[19] M. J. Wainwright. Information-theoretic limits on sparsity recovery in the high-dimensional and noisy setting. *IEEE Trans. Inf. Theory*, 55(12):5728–5741, 2009.

[20] A. K. Fletcher, S. Rangan, and V. K. Goyal. Necessary and sufficient conditions for sparsity pattern recovery. *IEEE Trans. Inf. Theory*, 55(12):5758–5772, 2009.

[21] S. Aeron, M. Zhao, and V. Saligrama. Information theoretic bounds for compressed sensing. *IEEE Trans. Inf. Theory*, 56(10):5111–5130, Oct. 2010.

[22] M. Akcakaya and V. Tarokh. Shannon-theoretic limits on noisy compressive sampling. *IEEE Trans. Inf. Theory*, 56(1):492–504, Jan. 2010.

[23] Y. Wu and S. Verdu. Optimal phase transitions in compressed sensing. *IEEE Trans. Inf. Theory*, 58(10):6241–6263, Oct.

[24] G. Reeves and M. Gastpar. The sampling rate-distortion tradeoff for sparsity pattern recovery in compressed sensing. *IEEE Trans. Inf. Theory*, May 2012.

[25] E. J. Candès. The restricted isometry property and its implications for compressed sensing. *Comptes Rendus Mathematique*, 346(9):589–592, 2008.

[26] Y. Chen and C. Caramanis. Noisy and missing data regression: Distribution-oblivious support recovery. In *International Conference on Machine Learning*, 2013.

[27] J. Wright and Y. Ma. Dense error correction via ℓ_1-minimization. *IEEE Trans. Inf. Theory*, 56(7):3540–3560, 2010.

[28] P. Diaconis and D. Freedman. Finite exchangeable sequences. *The Annals of Probability*, pages 745–764, 1980.

[29] D. Aldous. Exchangeability and related topics. *École d'Été de Probabilités de Saint-Flour XIII1983*, pages 1–198, 1985.

[30] S. L. Lauritzen, O. E. Barndorff-Nielsen, A. P. Dawid, P. Diaconis, and S. Johansen. Extreme point models in statistics [with discussion and reply]. *Scandinavian Journal of Statistics*, pages 65–91, 1984.

[31] R. G. Gallager. *Information Theory and Reliable Communication*. John Wiley & Sons, Inc., New York, NY, USA, 1968.

[32] E. J. Candès, M. B. Wakin, and S. P. Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14(5-6):877–905, 2008.

[33] R. G. Gallager. Information theory. In *Mathematics of Physics and Chemistry, Vol. 2*. Van Nostrand, Princeton, NJ, USA, 1964.